Author:
Sandra Bronsvoort

Supervisors:
University of Twente:
Dr. M.C. van der Heijden
Dr. E. Topan

Royal Schiphol Group:
Steven Kempen
Rob Sentveld

Long-term replacement planning for Royal Schiphol Group

An Integer Linear Programming model

Master Thesis, October 2018
Industrial Engineering and Management
Track Production and Logistic Management
University of Twente

Management summary

This research is conducted at the Asset Management division ("ASM") of Royal Schiphol Group. The fast growth of Amsterdam Airport Schiphol ("AAS") in the previous years has capacity in the terminal become scarce, demanding a different approach towards the planning of major maintenance activities and replacements. Till now on, decisions are made on asset-level and the vast size of the asset base at makes this approach very time-consuming and inefficient. The short planning horizon results in a low predictability, which in turn results in a low realization of plans, since it is hard to execute major maintenance activities while not disturbing the operational processes and the passengers. This arguably results in higher maintenance costs, since it happens that assets are kept up and running long after their economic end-of-life, by performing regular maintenance instead of replacing the assets. ASM believes that planning over longer horizons, as well as adopting a more integrated approach that combines the replacements of different asset types can increase the predictability and realization of plans and limit the impact on operations.

The main research question is therefore stated as follows:

> *How can the Asset Management division of Royal Schiphol Group plan the replacements of assets over a 60 year horizon, in order to limit the impact on operations?*

In order to develop a proof of concept, we focus in this research on replacements at the E-pier at AAS. To analyze the current situation at this specific pier, we want to estimate how the current approach, which plans over a 5-year horizon, would behave over a 60-year horizon. Since the realization of plans is currently low, this is hard to predict. We therefore referred to what we called the 'baseline situation', which is the situation in which we consider all assets individually and replace them immediately at the end of their economic life. We found that this practice, which is similar to the current approach, would result in replacements to take place in 51 of the 60 years. ASM recognizes that clustering some of these replacements may result in a planning that is more beneficial for the area as a whole and limits the impact on operations.

Clustering implies that assets are replaced earlier or later than their end-of-life, in order to combine their replacement with the replacement of other assets. This deviation from and asset's end-of-life may result in the individual asset not being optimally utilized and therefore comes at a penalty cost. Replacing an asset earlier than it's end-of-life represents a disinvestment costs, whereas postponing the replacement of an asset may result in higher maintenance costs and increased risks of failure. This results in a trade-off in the penalty costs and the number of clusters, which we both want to minimize.

In order to find the optimal planning that minimizes the number of clusters for the minimum costs, an Integer Linear Programming model was formulated and programmed in AIMMS. It is assumed that for each asset we know its lifecycle, the allowed number of years with which the asset is allowed to be replaced earlier or later than at the end of this lifecycle and the penalty costs for early or late replacement. The model is formulated such that it plans the repetitive replacements of assets over a 60-year horizon, while ensuring that the asset is not planned earlier than is allowed by the minimum lifecycle or later than is allowed by the maximum lifecycle. Early and late replacements are penalized. We cannot directly compare costs, since the model compares the penalty costs with the number of clusters, i.e. the number of years in which at least one replacement is planned. We therefore make use of a balancing parameter, that balances the importance of the penalty costs and the importance of the number of clusters. By changing the value of the balancing parameter, the decision maker is able to steer the model and accept more or less penalty costs.

When accepting more penalty costs, the number of clusters can be reduced more. A decrease in the number of clusters from 51 to 25 can for example be reached when ASM accepts in total €226,802 in penalty costs over the 60-year horizon. A decrease from 51 of 12 can be reached when Schiphol accepts €1,049,135 in penalty costs. An important conclusion is that the model chooses to align the replacements of assets as soon as possible by shifting the assets' first replacements, such that the subsequent replacements in the horizon are naturally in cadence and do not need to be shifted anymore. This way, a huge decrease in the number of clusters can be achieved by making relatively small shifts in replacement moments. This is important since it shows us that the impact on operations can be decreased without having to accept increased risks because of postponing assets with many years.

A sensitivity analysis was performed to see how the values for the input parameters influence the solution. We concluded that limiting the years with which an asset may be replaced early or late to only one year influences the solution planning and increases the penalty costs with 64.7% (when ASM wants to decrease the number of clusters to 25) and 190.6% (when ASM wants to decrease the number of clusters to 16). This happens because this setting hinders the model from early synchronizing the replacements and more shifts need to be made. This again stresses the importance of early alignment of replacement cycles.

Another important part of the sensitivity analysis challenges the assumption that the costs for late replacement of assets are linear in the number of years with which an asset was postponed. Moreover, in the original case study the costs for one year of late replacement were set to be more expensive than the costs for one year of early replacement. This resulted in many assets being replaced earlier than optimal and only a little number of replacements being postponed. In the sensitivity analysis we proposed an increasing, non-linear cost structure which in our opinion may better represent the actual situation. We now see that more replacements are postponed and relatively little assets are replaced earlier than optimal. The resulting penalty costs are much lower, but we advise ASM to invest in refining the cost functions in order to obtain a more truthful representation of the penalty costs.

The model is a helpful tool to ASM in the new strategy. In this new strategy the main contractor will have a more autonomous role, whereas the role of ASM will be more controlling. In this strategy, ASM provides the main contractor with time windows in which renewals and replacements of assets are planned. This model can guide ASM in identifying the optimal moments for these moments to take place. Moreover, planning over a longer horizon offers ASM more time to prepare for the clusters to be executed, since these moments are now known well in advance. If enhances predictability and offers the possibility to a better integration of the works in the operational processes. ASM expects that by clustering activities in less frequent moments the importance of these moments can be better communicated to internal and external customers and that these clustered moments offer a stage for developing and carrying out improvement projects as opposed to merely replacing individual assets.

Preface

This thesis was written in order for completion of the master programme Industrial Engineering and Management. I would like to thank Royal Schiphol Group for the opportunity to conduct my research here and especially my supervisors Rob Sentveld and Steven Kempen for their guidance and the great collaboration. I would also like to thank my other colleagues of the Technical Management team for all I have learned during this period and for the great time I had at Schiphol. Furthermore, I would like to thank my first supervisor Matthieu van der Heijden for his time investment, always valuable feedback and the constructive meetings. I would also like to thank Engin Topan who became my second supervisor in the last stages of my project and especially helped me gain new insights when I encountered some difficulties in the development of my model.

Table of contents

Management summary ... 3

Preface ... 5

Glossary ... 8

1. Introduction ... 9

1.1. Royal Schiphol Group ... 9

1.2. Problem statement ... 9

1.3. Research goal ... 11

2. Situation analysis ... 13

2.1. Developments and context ... 13

2.2. Current methodology ... 13

2.2.1. Description of the current situation ... 13

2.2.2. Clustering in the current methodology ... 16

2.3. Case: the E-pier ... 17

2.3.1. Performance of the baseline situation ... 17

2.4. Conclusion ... 18

3. Literature review ... 20

3.1. Preventive maintenance: benefits and disadvantages ... 20

3.2. Categorization of maintenance models ... 20

3.3. Exact models for maintenance clustering ... 21

3.4. Applicability to the case ... 23

4. Model ... 25

4.1. Model description ... 25

4.2. Mathematical formulation ... 25

4.3. Assumptions and conditions to ensure validity of the model ... 29

5. Case study ... 31

5.1. Input ... 31

5.2. Results for different values of the balancing parameter ... 33

5.3. Computing time ... 39

5.4. Sensitivity analysis ... 40

5.4.1. Allowed early (AEa) and allowed late (ALa) ... 40

5.4.2. Already fully depreciated assets ($Fa < 2019$) ... 41

5.4.3. Replacement value for constructive assets ... 42

5.4.4. Balance $CEa$and CLa ... 42

5.4.5. Non-linear penalty costs for late replacements ... 43

5.5. Conclusion ... 45

6. Implementation 47
7. Conclusion and recommendations 49
7.1. Conclusion 49
7.2. Recommendations 49
8. References 51
Appendix A: Asset database E-pier 52
Appendix B: Model formulation in AIMMS 53
Appendix C: Estimating replacements values 56
Appendix D: Sensitivity analysis, experiment 3 57
Appendix E: Modified model with non-linear costs for late replacement 58

Glossary

ASM	The Asset Management division of Royal Schiphol Group.
AAS	Amsterdam Airport Schiphol
Asset	Every individual unit that has a significant share in the total cost price of a system and is depreciated separately. All elevators are unique assets.
Asset type	All unique assets that fulfil the same function. For example: elevators.
Asset group	A set of assets that share the same asset type and construction year. For example: elevators built in 2003.
System	A set of assets that together deliver a value to the customer and fulfil a specified function, e.g. the passenger transport system, consisting of the asset types elevators, escalators and moving walkways.
Cluster	A set of replacements planned in the same year.

1. Introduction

As one of Europe's main airport operators, Royal Schiphol Group acts in a rapidly growing market with a rising demand for air transport. Aiming to efficiently fulfil this demand, the company faces complex decisions. In this chapter, we will introduce the problem on which this research will focus. Section 1.1 will provide more information about Royal Schiphol Group. In Section 1.2 we will elaborate on the problem context. In Section 1.3 the goal of this research, together with the associated research questions will be outlined.

1.1. Royal Schiphol Group

Royal Schiphol Group ("Schiphol") is an operator of airports and is the owner of Amsterdam Airport Schiphol ("AAS"), Rotterdam The Hague Airport and Lelystad Airport. It also has a majority share in Eindhoven Airport. Moreover, the company closely works together with foreign airports. The exploitation of AAS is the company's main activity. This thesis will focus on AAS only.

The activities at AAS can be subdivided into three business areas, i.e. Aviation, Consumer Products & Services and Real Estate. The key business area is Aviation, which provides service to passengers, airlines, freight handlers and logistics companies. Aviation develops and manages infrastructure that allow for an efficient and reliable movement of passengers, luggage and goods.

The Asset Management division ("ASM") is responsible for the planning, development, realization, management and maintenance of the approximately 45,000 assets at AAS. Examples of these assets can be the runways, passenger bridges, aircraft stands, but also climate systems, elevators or lighting.

The ASM division is divided into five subdivisions. The Strategy Office is responsible for aligning the ASM strategy for Aviation with the overall Schiphol strategy. Planning & Portfolio Management is responsible for translating the customer's demands into asset planning and Development is responsible for the actual realization of the asset. After realization the assets are transferred to Maintenance & Operations, which is the division that is responsible for the execution of maintenance on the asset during its life cycle. The fifth subdivision is the Technical Expert Center ("TEC").

TEC can be seen as the knowledge center of ASM. TEC draws, manages and improves asset policies and maintenance concepts, taking into account availability and costs, but also aspects such as legislation, sustainability and safety. TEC provides advise in order to optimize asset efficiency, steering at lowering costs, while ensuring that asset performance meets the standards as has been agreed upon internally and with customers. TEC can be further subdivided into four divisions, one of which is Technical Management ("TM"), the division in which this research will be carried out.

1.2. Problem statement

One of the main tasks of TM is the development of the so-called *meerjarenonderhoudsplan* (multi-year maintenance plan, "MJOP"). The main goal of the MJOP is to plan and budget major maintenance tasks, i.e. renovations and replacements, for the coming five years. The MJOP is updated every year based upon actual asset conditions and performances, resulting in a plan with a rolling horizon. In developing this MJOP, TM closely works together with Maintenance & Operations ("M&O") and the main contractors, who are responsible for performing the actual maintenance activities on the assets. When the main contractors and M&O feel that there is a need for maintenance on or replacement of an asset, a request is submitted to TM. TM evaluates the usefulness and necessity of performing maintenance or replacing assets. Risks and impacts are classified and based upon its priority, the maintenance task is scheduled somewhere in the next five years.

Under the current planning approach, decisions are made on asset level – or sometimes even on component level. This approach ensures that the assets are optimally utilized, but the extensive asset base also makes this approach time-intensive and complex. It also results in a high dispersion of activities over time and in many small projects being performed simultaneously. Since these projects have their own project teams, ASM also thinks that overheads costs can significantly be reduced when activities are more clustered. Activities in the MJOP, i.e. in the coming five years, are already clustered as much as possible in order to achieve economies of scale and lower the impact on operations. ASM however believes that planning over a longer time horizon increases predictability and therefore allows for easier integration in daily operations while limiting the disturbance for processes and passengers. The vast number of assets however makes it very hard to determine what optimal packages of maintenance tasks and replacements and when to carry out these clusters.

In addition, ASM believes that the new approach will increase the (timely) realization of maintenance plans. Plans are now often postponed, since maintenance and replacements almost always interfere with the daily operational processes at the airport and may therefore decrease the passenger's comfort. ASM thinks that planning longer in advance makes it easier to integrate maintenance in the daily process, since there is more time to come up with additional measures to limit the inconvenience for the passenger. Another reason for the postponements is that maintenance turns out to be hard to sell to customers. It is often not clear to airlines what the added value of maintenance is. This often results in maintenance projects being postponed in order to free budget for new developments.

The fact that ASM is currently on the verge of an organizational change is important for understanding the context of these problems. Schiphol's strategy is to operate in accordance to a control model. This allows Schiphol to focus on its core processes and to make optimal use of the expertise of its business partners and suppliers. In the current situation, the actual execution of maintenance is already outsourced to the main contractors. ASM is however still highly involved in developing maintenance concepts, evaluating the need for maintenance and replacements and scheduling on operational level. In the new methodology, the main contractor will have a higher responsibility and act more autonomous. ASM will provide the main contractors with a long-term planning which defines the moments at which major maintenance, overhauls and replacements of assets are planned. This way, ASM is going to make a shift from result-based to performance-based contracts. The responsibility of the main contractor would be to ensure that the asset, or a process as a whole, meets the predefined performance levels until these moments of intervention, by performing regular maintenance. Also the management of asset data, and making predictions on asset performance, degradation and failure behavior on asset level will become the responsibility of the main contractor. In addition, till now contractors were responsible for one technical discipline, for example the buildings itself, building-specific installations and energy production, fire safety or operating assets. From 2019 on, the contractors will be responsible for a geographical area with all technical disciplines within it. The coordination between the different technical disciplines and individual assets therefore becomes more important.

In summary, ASM wants to make a shift from the current bottom-up approach in which decisions are made on asset level to a top-down approach in which maintenance plans are made by considering areas of the airport and their processes as a whole. By planning over a longer time horizon, ASM wants to increase predictability and the level of realization. Moreover, a long-term, integrated strategy offers more opportunities for clustering to lower the impact on operations and reduce costs. This research focuses on the development of a model that plan these replacements over a longer planning horizon.

1.3. Research goal

ASM decided that for now the focus should be on the long-term planning of asset replacements. Because of time constraints, the focus of this research will be on the E-pier. This way, a proof of concept will be delivered, which can later on be extended to other areas at AAS. The lifetime for the construction of a pier is 60 years, which is why a planning horizon of 60 years is taken. Many different stakeholders are involved in replacement activities in the terminal. The focus of this research will however be on developing a planning that is optimal from an ASM viewpoint. Therefore, the main research question will be as follows:

> *How can the Asset Management division of Royal Schiphol Group plan the replacements of assets over a 60 year horizon, in order to limit the impact on operations?*

The research question will be answered by answering the following sub questions:

1. What is the current situation and how does the current methodology perform?
    a. What are relevant developments in the aviation industry? How do these developments influence the context at AAS?
    b. What is the current situation? How does this methodology perform? Based on what characteristics are activities clustered at the moment?
    c. What does the asset base of the E-pier look like?

    In order to gain insight in the current situation, in Chapter 2 we will discuss the relevant developments in the aviation industry and the context at AAS. We will analyze the current situation and its performance to see what problems result from the current methodology and if there is indeed potential for improvement. Furthermore we will zoom in on our case: the E-pier.

2. What is written in existing literature about maintenance optimization?
    a. What models for maintenance optimization are known?
    b. For what situations are these models suitable?
    c. What are the strengths and weaknesses of these models?
    d. What methods are relevant for ASM?

    In Chapter 3, existing methodologies for maintenance optimization are reviewed. Based on their characteristics, we will examine which methods can act as a basis for a model that suits ASM.

3. How can ASM optimally plan its replacement activities in the E-pier?
    a. How can we determine the deviation from the optimal moment for an individual asset?
    b. How can we express the costs that arise from this deviation?
    c. What constraints should be incorporated in the model?
    d. What assumptions and conditions have to be met in order to ensure validity of the model?

    In Chapter 4, the techniques found in Chapter 3 are used to develop an optimization model that fits the specific context in which ASM operates. This model should balance the costs and benefits of clustering replacement activities. We should determine how we express these costs and benefits.

4. What are the benefits of the new methodology?
    a. How do the different input parameters influence the outcome of the model?
    b. What savings can the new methodology obtain? How much would ASM have to invest to achieve these savings?

    In Chapter 5 we will analyze the benefits of long-term replacement planning and the clustering of replacement activities. We will research how the various input parameters influence the solution.

5. How should the new methodology be implemented?
    a. What practices do we recommend for the use of the new methodology?
    b. What conclusions can we draw from this research?
    c. On what areas should be focused when ASM wants to develop the model further?

    In Chapter 6 we will give an advice regarding the implementation of the model. In Chapter 7 we will elaborate on our conclusions and recommendations.

2. Situation analysis

This chapter answers the first sub question: *'What is the current situation and how does the current methodology perform?'.* Section 2.1 will briefly discuss recent developments in the aviation industry and how these influence the context at Schiphol. Section 2.2 will provide insight in the current methodology. Lastly, Section 2.3 will assess the performance of the current methodology for our case, the E-pier, in specific.

2.1. Developments and context

Over time, AAS has become one of the best connected hub airports in Europe. At the moment, the airport has 326 direct destinations. A wide range of factors such as economic recovery, growing world trade, low oil prices and a higher competition between airlines, has led to a rapid growth of the aviation industry over the last years (Royal Schiphol Group, 2017a). 2017 was AAS's busiest year ever with 68.5 million passengers: a growth of almost 8% with respect to the year before. The number of seats per air transport movement has increased from 165 in 2016 to 168.6 in 2017. Simultaneously, the average passenger load factor has increased from 83.8% in 2016 to 84.7% in 2017 (Royal Schiphol Group, 2017b). This means airlines not only fly with larger aircraft, but that also the number of seats is used more efficiently. Although AAS has almost reached its air transport movement ceiling of 500,000 starts and landings per year, the number of passengers is therefore expected to grow even further.

The consequences of this growth are twofold. First, increasing passenger numbers imply that the load imposed on the assets at the terminal complex increases too. This might result in undercapacity of for example air treatment- or cooling systems and might accelerate the degradation process of these assets. The need for maintenance might therefore increase and assets may have to be replaced earlier than initially estimated.

At the same time, availability and capacity of the terminal becomes increasingly critical in order to be able to house all these passengers and handle the boarding-, transfer and security processes and smooth flow of passengers. This makes it even more difficult to conveniently integrate major maintenance works and replacements in the daily operational processes. For stakeholders like Schiphol's Operations department, Security or airlines, the main priority is continuity and a minimal disturbance of the daily processes. In earlier years, when there was more flexibility in the capacity at AAS, the short-term approach was sufficient since maintenance could be planned more easily alongside daily processes. Nowadays however, predictability of ASM's plans becomes more and more important, in order for the different stakeholders in the terminal to be able to prepare for maintenance activities that will heavily impact the processes. In addition, ASM's experience with major maintenance at airside, i.e. for example at the runways, has shown that many stakeholders prefer longer, but less frequent disturbances over more dispersed, but smaller disturbances. This demands planning over longer horizons.

2.2. Current methodology

In Section 2.2.1 we elaborate on the current situation and we will analyse to what extent major maintenance tasks and replacements are indeed postponed and what the effects of these postponements are. In Section 2.2.2 it is described how and to what extent major maintenance tasks and replacements are currently clustered.

2.2.1. Description of the current situation

Maintenance concepts have been developed for every type of asset at Schiphol. In these maintenance concepts, the optimal balance between preventive and corrective maintenance has been determined. Moreover, it is stated what maintenance tasks should be executed when and with what frequency. The maintenance concepts often only include regular maintenance, which are the relatively small maintenance tasks that can be executed without significantly disturbing the operational processes and have a repetitive character. Examples can be the weekly cleaning, a monthly inspection, yearly lubrication or the replacement of small components. Additional corrective actions are then taken in case of malfunctioning or failure of the asset. These activities do not have their own budgets, but are categorized as operational expenses. On the contrary, major maintenance activities like midlife upgrades, overhauls, refurbishments and replacements do have their own budgets. They aim at significantly improving the current state of the asset and are therefore classified as investments. The duration of these projects is relatively long and they are likely to impact daily operations. As indicated, the planning of these major maintenance tasks and replacements is done in the MJOP.

Decisions for heavy maintenance and replacements are condition-based. In the current methodology, assets are monitored individually, resulting in decisions for replacements being made on asset level. Replacement years can be determined by looking at the expected useful life of an asset, which is defined to be the shortest of the expected technical lifetime and the expected economic lifetime of the asset (Royal Schiphol Group, 2017a). The technical lifetime of an asset is the duration that the asset is functional. The economic end-of-life is the moment from which it becomes financially more attractive to replace the asset for a new asset, for example because the asset becomes less reliable and maintenance costs increase or a new asset can fulfil the desired function more efficient. Determining the useful life of an asset is outside the scope of this research and we assume that either the technical- or economic end-of-life – whichever comes first – is indeed the optimal moment for replacement. Since decisions are made on asset level, it would theoretically be optimal to replace assets at the end of their economic end-of-life. In practice, we see that replacement is often postponed.

Most of the activities on the 2019 MJOP did already appear on the 2018 MJOP. Although many of these activities where classified as activities with a high priority, they again appear on the MJOP of 2019, meaning that they have not been executed yet. Relatively little of them will actually start in 2019. There may be additional activities on the 2018 MJOP that were planned for execution, but may in reality not have been executed. The enormous extent of the MJOP however makes it however difficult to follow-up all activities. There is no general database that keeps track of the realization of the MJOP. We do not know to what extent the MJOP plans have been realized and actual percentages of postponed activities may therefore be even higher. What we do know is that around 50% of the budget of the 2018 MJOP was realized. This does however not say much about the number of activities that was actually executed. It might for example be the case that more than 50% of the activities was executed, but that the related costs were much lower than foreseen.

There are multiple valid reasons for postponing MJOP activities, for example to combine them with future renovation projects to save costs. Many activities are however postponed after finalizing the MJOP, since the scarce capacity in the terminal does not leave any room for major maintenance and replacements to take place. Assets are then kept up and running by performing regular maintenance.

For this research, the most actual asset database was used, containing information about all active assets at AAS, i.e. the assets that are currently in use. The database consists of over 45,000 assets. For our research only the assets in plot 5, related to the terminal and piers are relevant. Filtering the data on location, we find that around 22,393 assets are located in these areas.

Examining the database confirms the tendency to postpone replacements. For the assets for which the depreciation period could be determined, it was found that 34% of these assets is fully depreciated. Based on their economic end-of-life, one would therefore expect that these assets would already have been replaced. For 20% of the asset base, the economic end-of-life was exceeded with five years or more. In Figure 1 the economic end-of-lifes of the active asset base are plotted in a map of AAS. In red the percentage of assets for which the economic end-of-life was more than five years ago (< 2014), in orange the percentage of assets for which the economic end-of-life was reached somewhere in the last five years (2014 – 2018) and in green the percentage of assets that has not reached its economic end-of-life yet (> 2019). A more detailed view for the E-pier in specific is displayed in Figure 2.

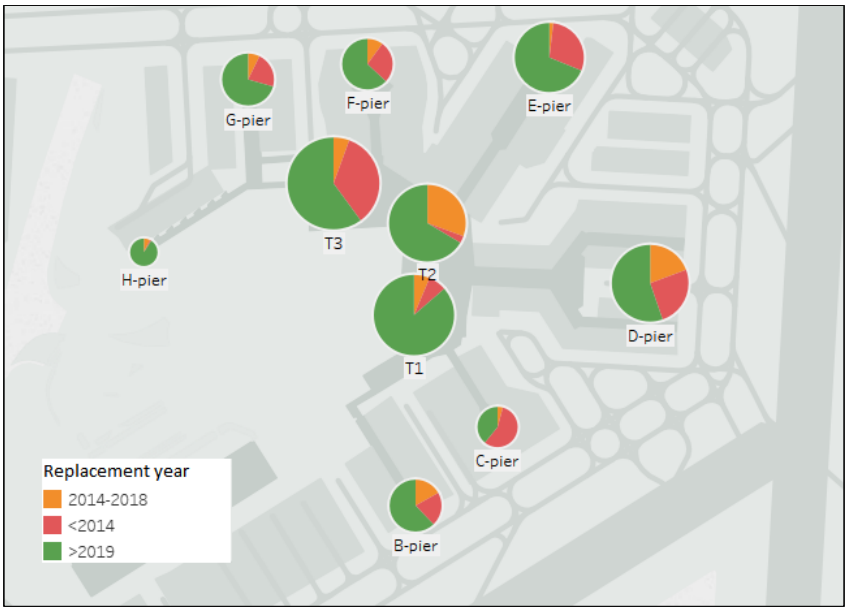

Figure 1. The economic end-of-life of the active assets represented in three categories and visualized in pie-charts over a map of AAS.

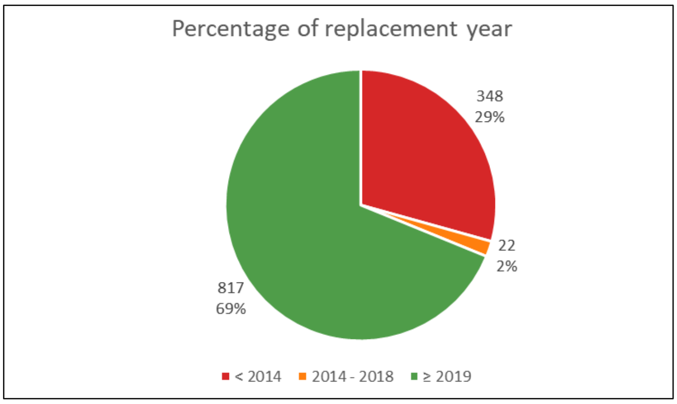

Figure 2. The percentage of active assets in the E-pier that has already reached their economic end-of-life is 31%. It can be seen that of these assets, most exceeded the economic end-of-life by 5 years or more.

It is important to realize that it is not necessarily a bad thing when assets are already fully depreciated. A probably better conclusion is that the depreciation period was determined on conservative estimations of the assets' economic life spans, which may be perfectly justified for financial reasons. In addition, it is expected that there are gaps between assets' economic and technical life span, since the assets are technically still in good condition. However, when the depreciation period correctly represents the economic lifecycle of an asset, this end-of-life is indeed the optimal moment to replace an asset. Postponing these replacements by patching-up the assets is expected to result in an increasing need and costs for maintenance. The graph in Figure 3 seems to confirm that there is a relationship between the ageing of the asset base and the number of workorders. Along the x-axis the percentage of assets that has already exceeded its economic end-of-life is plotted. The y-axis displays the average number of workorders per asset in the period 2013-2017, which is calculated by dividing the total number of maintenance workorders at a certain location by the total number of assets at that location. These workorders can have both a corrective or a preventive character. The graph shows a positive correlation between the percentage of assets that was already fully depreciated and the number of work orders. It is important to notice however, that many other factors may influence this relationship. It might for example be the case that assets at Terminal 1 (T1) are used more extensively than assets at the H-pier. Moreover, it can be argued that other measures are more meaningful than the number of workorders, since this measure does for example not take into account the duration of a workorder or the costs related to it. Because of data availability and reliability, we will however not further investigate these relationships.

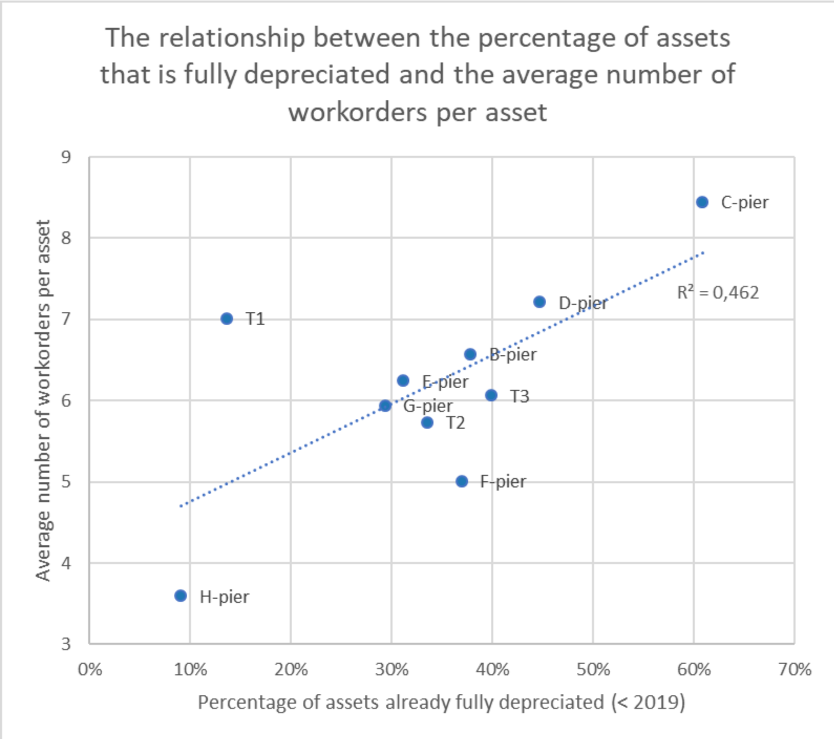

Figure 3. The average number of workorders in the period 2013-2017 per asset per location plotted against the percentage of assets at that location that have already exceeded their economic end-of-life.

Besides higher maintenance costs, keeping assets in operation much longer than their economic lifetime may have other serious consequences. Assets may be still functioning satisfactory long after their economic end-of-life, but this may come at a risk. Failure of the asset may then result in both high costs and operational disruptions, because spare parts have become obsolete and are not available anymore. Situations like this have not taken place yet and ASM closely monitors the condition of each asset, so this is not probable to happen in the near future either. When the depreciation period however correctly represents the economic lifetime of the asset, it is financially more beneficial to replace the asset for a new model instead of keeping maintaining the old asset. In addition, for many asset types it is likely that a new asset is more efficient than the old asset, for example in terms of output or energy usage.

2.2.2. Clustering in the current methodology

Maintenance tasks and replacements in the MJOP, i.e. in the coming five years, are at the moment clustered based on technical function and location. As an example, the replacement of four air treatment units in the same technical room may be combined into a single cluster. Replacing these four units altogether limits the impact on operations and might reduce set-up costs or costs for necessary equipment. There is however still room for improvement. When ASM starts to plan over longer time horizons, activities are in sight longer in advance, resulting in more possibilities for clustering. In addition, in the new methodology there will be a higher level of integration between the different technical disciplines since the main contractor will be responsible for a geographical area as a whole. This also allows for more possibilities in the combination of various activities.

Although the new approach offers more possibilities in clustering, the large size of the asset base and the MJOP makes it hard to manually cluster activities. Before clustering, the MJOP has a size of more than 5,000 rows, where each row corresponds to a maintenance task or replacement. It is not possible, nor desirable, to manually assess all possibilities for clustering. The grouping of activities is at the moment therefore rather pragmatic, instead of a standardized data-based decision-making process in establishing the MJOP.

2.3. Case: the E-pier

For this research, the E-pier at AAS will be taken as a case study. Being built in 1987, the E-pier is one of the oldest piers at AAS. Besides one narrow-body stand, all other 13 stands are suitable for the handling of wide-body airplanes, mainly used by KLM. The E-pier consist of four levels, i.e. the basement, the ground floor, the first floor and the second floor. The basement houses part of AAS's luggage handling system, whereas the ground floor houses offices of different airlines. The first- and second floors are the passenger areas. The second floor has been build more recently, i.e. in 2015, as part of the One-XS program in which Schiphol switched from decentralized to centralized security filters.

Taking the E-pier as a case study means that we will focus on gathering data for the assets in this pier and that the model will be tested on the E-pier's asset base. It is important that we test the model with a proper representation of the total asset base of the E-pier in order to be able to assess the performance of the current situation and get a good idea of the performance of the model later on. 1,034 were included in the case. An overview of these assets can be found in Appendix A. In Section 2.3.1 we will elaborate on the performance of the current methodology for this dataset.

2.3.1. Performance of the baseline situation

In order to assess the performance of the current methodology over the long term, we have to estimate how we expect the current methodology, which plans over a five-year horizon, behaves over a time period of 60 years. This is difficult since the realization of the plans is currently relatively low. We will therefore differentiate between the current methodology and the 'baseline situation' from now on. The baseline situation reflects the situation in which assets are replaced immediately at their end-of-life. Similar to the current methodology, in the baseline situation replacements are planned for all assets separately. In the baseline situation, clustering is therefore merely an indirect result of coinciding replacement moments, but no well-considered decision intending to reduce the impact on operations.

As mentioned in Section 1.3, ASM wants to focus on the replacement of assets. The main goal of developing a long-term replacement planning is to decrease the impact on the operational processes. ASM believes that the impact on operations can be reduced by clustering the many individual replacements of assets in larger, but less dispersed projects. As described, in the baseline situation assets are replaced at their economic end-of-life, without taking into account the possible benefits of clustering. When we plot the economic end-of-life of the E-pier assets over time, we can conclude that the replacements are highly scattered over time. In 51 out of 60 years at least one asset should be replaced. This is visualized in Figure 4. For the sake of clarity, the asset types are grouped into systems, where a system is defined as a group of assets that together fulfil a specified function, e.g. the climate system. The colours refer to the different systems. In 2040 for example, at least one asset in the systems climate, low voltage, elevators and escalators reach their economic end-of-life and should therefore ideally be replaced.

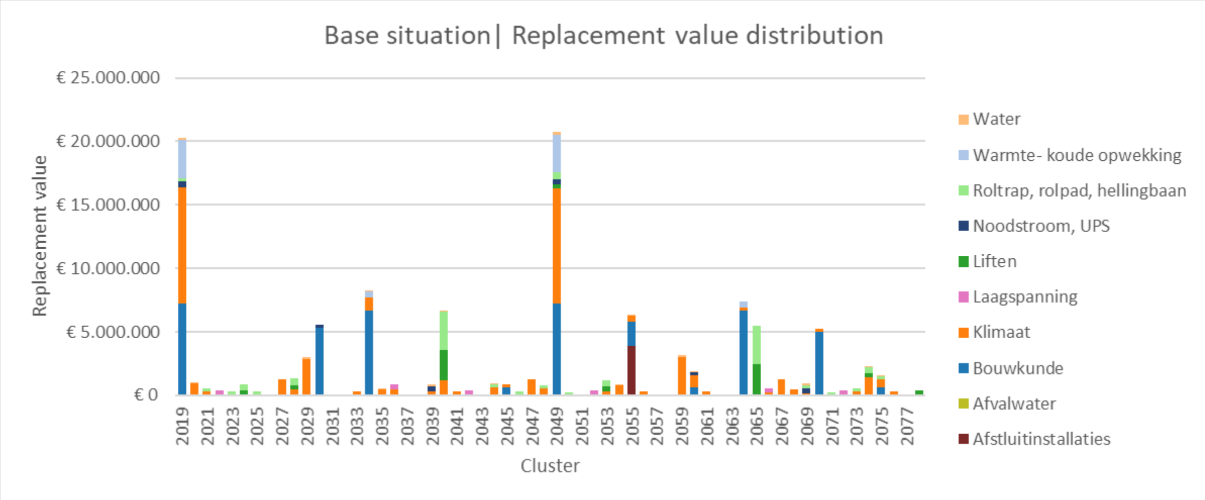

Figure 4. The replacement moments of the assets in the E-pier over the time horizon, resulting in a high number of clusters (51).

ASM believes such a high dispersion of replacements over time comes with many disadvantages. Often mentioned by the different experts that were interviewed is the fact that almost every single replacement is approached as a separate project with its own administrative-, engineering -, and project costs. It is often heard that the lack of clustering results in preparatory activities being executed multiple times. A well-known example is opening up the ceiling in a certain area for the replacement of lighting, closing the ceiling, and opening it again a year later because the sprinklers have to be replaced too. When these replacements were executed together, the needed tools, equipment and labor could have been shared. Maybe even more important is the impact on capacity. During large replacements, parts of the terminal – or in our case the E-pier in specific – may be inaccessible to passengers, resulting in a disturbed passenger flow or capacity losses. ASM believes that clustering replacements in less but heavier moments, as opposed to many smaller moments highly scattered over time, can decrease the disturbance of the operational processes and result in significant financial benefits. Moreover, ASM thinks the development of a long-term replacement planning has the ability to improve the passenger perception.

Clustering however comes at a cost. We assume that for deviating from the optimal individual replacement moments penalty costs have to be paid. Early replacement of an asset represents a disinvestment, whereas postponed replacement leads to additional maintenance costs. Since in the baseline situation assets are replaced immediately at their economic end-of-life and there are no penalty costs resulting from early replacement or postponement. Therefore, in the baseline situation the penalty costs are €0. How much penalty costs ASM wants to accept to decrease the number of clusters may differ from location to location, since the preferred outcome of the model may depend on the preferences of internal and external stakeholders at a location. Although a high dispersion is often undesirable, it also results in smaller work packages per cluster and a lower average duration per cluster. In some situations this may be preferred over a low number of clusters. For ASM it is therefore important that the model to be developed can be steered towards more or less clustering in order to be able to roll-out the model to other locations besides the E-pier and take into account the preferences of different stakeholders.

2.4. Conclusion

In this chapter we have answered the research question *'What is the current situation and how does the current methodology perform?'*. We can answer this question by concluding the following five things:

- *Developments in the aviation industry demand a different approach towards the planning of replacements.*
  The aviation industry is growing and every year more passengers visit Schiphol. This implies that assets are used more intensively and may increase the need for maintenance. At the same time, capacity becomes more scarce, which demands for ASM to plan over longer time horizons.

- *Many assets are already fully depreciated.*
  This is most likely caused by a conservative estimation of the life spans of the assets, which is not per se a bad thing. When the depreciation period however correctly reflects the economic lifetime of the asset, it may be the case that replacing the old asset for a newer one would financially be the better choice.

- *The new methodology offers potential for more clustering.*
  Since in the new methodology ASM will plan over a longer time horizon and the main contractors will now be responsible for all technical disciplines in a geographical area, there are more possibilities for clustering. The vast size of the asset base makes it however impossible to manually assess all options.

- *The baseline situation results in a high dispersion of activities.*
  - When assets are replaced immediately at their end-of-life, without considering clustering, ASM would have to replace assets in 51 of the 60 years.
  - The penalty costs in this situation are €0, since replacements are not shifted.

- *What is optimal may differ per location and stakeholder.*
  ASM should therefore be able to steer the model in a preferred direction, putting more weight on the costs, or more weight on the number of clusters.

3. Literature review

This chapter aims to answer the second sub question: *'What is written in existing literature about maintenance optimization?'*. We will start this review with brief overview of various maintenance types in Section 3.1. In Section 3.2 we will focus on common categorizations in various types of maintenance optimization models. Section 3.3 discusses several exact models on the clustering of maintenance activities. Lastly, in Section 3.4 we will the applicability of these models to the situation at Schiphol.

3.1. Preventive maintenance: benefits and disadvantages

Maintenance is defined as those activities that are performed in order to retain systems in, or restore systems to the state that is necessary for fulfilment of its function (Gits, 1992). Corrective maintenance is carried out after a system has failed and therefore reactive in nature. Costs for corrective maintenance are likely to be high, because failure of an asset might result in system downtime, safety dangers, or might cause additional damage to other assets. As opposed to corrective maintenance, preventive maintenance is performed in order to prevent the system from failing and thus has a more proactive character. A special type of preventive maintenance is condition-based maintenance. In condition-based maintenance the execution of maintenance is triggered by inspections or condition measurements (Budai-Balke, 2009). Many argue that this is more effective and efficient than preventive maintenance that is for example solely based on the age of assets. Budai-Balke (2009) however states that predicting failures is often very difficult, which makes it hard to plan maintenance in advance. Budai-Balke (2009) also states that for complex systems it is very hard to monitor all individual units, as well as organize all information in databases. In such complex systems, scheduled maintenance based on ageing might for example be more convenient. Although preventive maintenance aims to minimize the disadvantages of corrective maintenance, it can also result in additional costs since it is likely to result in more maintenance than is strictly needed.

3.2. Categorization of maintenance models

For maintenance models in general, as well for models on clustering in specific, three important categorizations can be recognized. First, the differentiation between single-component and multi-component models, second the differentiation between long and short planning horizons, and third the differentiation between deterministic and stochastic models.

The first differentiation, i.e. between single-component and multi-component models, is most straightforward. A single-component or single-unit model only considers one specific component, whereas multi-component models aim to optimize maintenance policies for a system consisting of several components with or without dependencies between them (Cho & Parlar, 1991).

Moreover, a distinction can be made between long-term and short-term maintenance models. Long-term planning for example focuses on the determination of execution moments of (major) maintenance activities or maintenance clusters that need to be aligned with other plans, whereas short-term scheduling for example deals with determining the order of execution activities, priority setting and the efficient use of the labor pool (Budai-Balke, 2009; Dekker, 1996; Van Dijkhuizen & Van Harten, 1997). In their review on maintenance models with economic dependence, Dekker, Wildeman and Van der Duyn Schouten (1997) explain that in long-term maintenance planning, it is often assumed that situations are stable over the long-term. These models often plan over infinite horizons and generate static planning rules that do not change over this horizon. Examples are models that generate long-term maintenance frequencies, meaning that activities will always be executed at the same time until the end of the horizon. Wildeman, Dekker and Smit (1997) state that infinite horizons are often applied as an approximation of a long-term stable situations, but that in reality planning horizons are usually finite for various reasons. Information is for example often only available for the short term and modifications of the systems may completely change the problem. In finite horizon models, the implicit assumption is made that the system is not used after the horizon. At the end of the horizon the system has totally lost its value or the system is worth its residual value (Dekker et al., 1997). However, assets often have longer lifetimes than the length of the horizon. As a result, many models use a rolling horizon approach (Budai-Balke, 2009; Wildeman et al., 1997). Rolling horizon models aim to capture the advantages of finite- and infinite horizon planning. These models plan over finite horizons, but decisions are based on long-term static planning over the infinite horizon. As Dekker et al. (1997) explain, when planning over rolling horizons, the preliminary long-term planning is adapted to the short-term situation. Decisions for the current finite horizon are implemented and afterwards a new horizon is considered. These models are dynamic in the way that they provide planning rules that can change over the planning horizon, by taking into account non-stationary events. Examples of such non-stationary are varying use and deterioration of assets or unexpected maintenance opportunities that allow for executing maintenance at lower costs.

A final distinction can be made between deterministic and stochastic models. A deterministic model is a model in which, for any value of the decision variables, the corresponding value of the objective function as well as whether or not the constraints will be satisfied is known with certainty (Winston, 2004). In stochastic models, this is uncertain. Likewise, for maintenance in specific, deterministic problems are defined as problems in which the timing and the outcome of maintenance actions are assumed to be certain, whereas in stochastic models this depends on chance (Budai-Balke, 2009). Within stochastic models, Dekker (1996) makes a further distinction between models under risk and models under uncertainty. Here, risk is described as the situation under which the probability distribution of the time to failure is known, whereas in case of uncertainty this distribution is unknown.

3.3. Exact models for maintenance clustering

An effective method in reducing maintenance costs can be the simultaneous execution of planned maintenance activities, which is often referred to as clustering (Budai-Balke, 2009). Clustering might be beneficial when there is some form of dependence between assets. If assets are dependent on each other, what is optimal for one asset is not necessarily optimal for the system as a whole (Cho & Parlar, 1991). Clustering individual maintenance- or replacement moments may therefore have benefits over the individual execution of these activities. At the same time, deviating from these single-asset optimal moments may come at a certain cost, for example because the clustering results in some activities having to be performed more often than originally planned (Budai-Balke, 2009). Many papers deal with the clustering of maintenance- and replacement activities, where the main goal is often cost reduction by combining activities to save on preparatory costs, such as downtime costs, needed equipment or the travelling of maintenance crew. These so-called set-up costs can be saved when activities are simultaneously executed, since only one set-up is required for the execution of a group of activities. Hence, the aim of the models that are developed in these papers is to find the optimum in the costs for deviating from the optimal moments for individual replacements and the benefits of combining these separate activities (Dekker, Smit & Losekoot, 1992; Van Dijkhuizen & Van Harten, 1997; Wildeman et al., 1997). The problem definition is often in accordance with the following formulation based on the problem in Wildeman et al. (1997). Consider a multi-component system with n components i. A preventive maintenance activity can be carried out at each component. These activities are assigned a component-dependent costs c^p_i, as well as a set-up cost S which is equal for all activities in the system. When an activity is executed separately, a cost of $c^p_i + S$ is incurred. For a group of activities that are executed simultaneously, only one time the set-up costs plus the sum of the component-dependent costs has to be paid.

When reviewing the existing literature, many different approaches to clustering can be found. Let us start with the relatively simple model as proposed in Liang (1985). In this model, the execution moments for individual activities is bounded by time windows. The problem is to find the optimal combination of activities by minimizing the sum of absolute deviations from the execution moments. The approach is pragmatic and helpful when little to no data is available, but it does not include a method to balance the costs of deviating from original execution moments with the savings resulting from the combination of activities. As a result, the model may output multiple different solutions, without being able to determine which one is the best. Moreover, it is assumed that early execution and late execution comes at the same costs and that these costs are equal for all maintenance activities.

Dekker et al. (1992) have extended the heuristic of Liang (1985) by integrating a cost component. In their version of the problem, deviation from the individually planned maintenance activities is not bounded, and one can therefore alter them. The only restriction is that each activity should be performed within the planning horizon. The problem is formulated as a set-partitioning problem, splitting up the set of all activities into subsets where the activities that together form a subset are executed simultaneously. The model aims to find the optimal partitioning, i.e. the partition that minimizes total costs. Since the number of set partitions grows exponentially in the number of maintenance activities complete enumeration of the combinations is impractical. The authors present some theorems that reduce the problem size. The authors acknowledge that in reality penalty functions are often difficult to obtain. In these situations they advise to multiply the absolute deviation from the originally planned execution moment by a scaling factor. If this factor is defined to be low enough, the model will maximize the number of combined activities and simultaneously minimize the sum of the deviations. Dekker et al. (1992) also recognize that identifying the savings resulting from the clustering of activities is in practice very difficult too. Therefore, they assume that all maintenance activities can be divided into groups that share the same preparative work that is unique for that group and that activities in different groups do not share the similar set-up costs. This simplifies the problem since one only has to consider combining activities within one group and that this combination results in the same savings.

In their paper, Wildeman et al. (1997) extend the model in Dekker et al. (1992) and formulated it as a dynamic programming model. They propose a rolling-horizon approach with five phases.

1. *Phase 1: decomposition.* In this first phase the frequency of the maintenance activity is optimized over an infinite horizon, resulting in maintenance rules for each separate activity. In this phase an average use and deterioration is assumed. Also interactions between components are neglected.
2. *Phase 2: penalty functions.* For each activity, the additional expected costs of deviating Δt from the execution time as determined in phase 1 has to be determined. These penalty functions are usually derived from the maintenance models in phase 1.
3. *Phase 3: tentative planning.* From this phase on, the planning horizon is considered to be finite. Based on the individual maintenance rules of phase 1, together with the current state of the component and additional short-term information, the time t_i at which the activity is carried out if it where independent of other activities is determined.
4. *Phase 4: grouping maintenance activities.* In this phase it is possible to deviate from the tentatively planned execution times for the individual components, in order to make it possible to execute them simultaneously. The optimal grouping structure maximizes the set-up cost reduction minus the costs of deviating from the tentative execution moments, i.e. the penalty costs.
5. *Phase 5: rolling horizon step.* After applying phase 4, the decision maker can manually change the planning and return to phase 3, which is an iterative process which can be done as often as desired. When the decision maker is satisfied, the grouping of phase 4 will be carried out and when planning for a new period, phases 3, 4 and 5 are repeated.

Another clustering model is the Preventive Maintenance Scheduling Problem (PMSP) in Budai-Balke (2009). This is a problem in railway maintenance in which short repetitive activities have to be combined with large projects over a finite horizon and in deterministic time slots. For each routine work, the maximum period between two consecutive executions is given. Moreover, it is known when the activity was executed most recently. It is possible to execute a maintenance activity earlier than necessary, i.e. earlier than the end of its interval. As a result, it is not known beforehand how many executions there will be in the planning horizon. Also a list of projects together with their duration and earliest and latest possible starting times is known. Too early or too late execution of projects is penalized with a cost. Set-up costs are here defined as track possession costs, which are mainly determined by the time that a certain track is unavailable for railway traffic because of maintenance works. The goal is to minimize the sum over the track possession costs and the maintenance costs. The problem is formulated as an linear model. The PMSP can be extended by fixing the intervals between two consecutive executions, i.e. earlier execution is not allowed anymore. This extended version of the PMSP is called the Restrictive Preventive Maintenance Scheduling Problem (RPMSP).

3.4. Applicability to the case

In Section 3.1 we have discussed several types of maintenance. The focus of this research will be on replacements only. As we have seen in Chapter 2, decisions regarding replacements are currently based on the conditions of the individual assets. As suggested by Budai-Balke (2009) this might have several disadvantages. The high complexity and enormous size of the asset base indeed result in a very time-consuming decision making process, losing track of planned and actually executed activities, a high number of postponements and a low level of realization. In line with Budai-Balke (2009), we therefore want to develop a planning in which replacements are planned long in advance.

We saw in literature that when there is dependency between assets, the optimal maintenance schedule for a single asset does not necessarily have to be optimal for the system as a whole as well. This is also the case for ASM. It aims to plan major maintenance on and replacement of assets based on the assets' individual conditions, since this is assumed to be optimal. However, when zooming out and considering the area – in this case the E-pier – as a whole, we see that this approach results in a high impact on operations caused by the high dispersion of activities. In other words, clustering activities by deviating from these individual execution moments might enable us to find a planning that better fits the preferences of ASM and its customers.

Our approach will follow the framework of Wildeman et al. (1997), where we will focus on phase 4, i.e. the clustering of activities. For this fourth phase we will develop a model based on the work of Budai-Balke (2009). We assume that the frequencies of activities are already optimized in phase 1 of the framework and that this has resulted in optimal replacement years based on the economic end-of-life. We assume an average use and deterioration, which results in the tentative planning in phase 3, which is what we called the baseline situation in Chapter 2. The model is therefore of a deterministic character, i.e. it is assumed that it is known in advance when replacements should be performed. We will also work with a finite horizon. It is therefore assumed that after the horizon the assets will not be used anymore and lose their value. There is no incentive for executing maintenance just before termination of the horizon in order to end up with 'healthy' assets. It is instead assumed to be beneficial to let assets degrade towards the end-of-the horizon, so that the pier can be renovated or demolished as a whole without having to disinvest newly placed assets. The model to be developed should form the basis for a rolling horizon approach, in which the long-term planning will be challenged based on short-term information, as in phase 5 of Wildeman et al. (1997). This step is however outside scope of this research, but is part of the implementation phase.

To tackle the problem in phase 4 of Wildeman et al. (1997) we will develop a model based on Budai-Balke (2009). There are although some significant differences in the problem definition. As we saw, the main goal in Budai-Balke (2009) is to minimize the sum over the maintenance costs, penalty costs and possession costs. In our case we do not directly take into account yearly maintenance costs. We assume that the asset should be replaced anyway, but that the moment of this replacement can be shifted forwards or backwards. In case of an early or late replacement, penalty costs have to be paid. In Wildeman et al. (1997) these costs are determined already in phase 2. In Budai-Balke (2009), penalty costs are charged for the situation in which the last execution is carried out too early compared to the end of the horizon, i.e. the last cycle is too long. We do not take into account such costs, since we assume that the assets should degrade towards the end of the horizon.

Moreover, in Budai-Balke (2009) intervals are restricted (in the RPMSP) or can be executed earlier but not later (in the PMSP). In our case, execution moments can be shifted in both directions, i.e. replacements can be planned both earlier and later than at the end of an asset's lifecycle. Moreover, shifts are bounded by a maximum number of years. In addition, our definition of the possession costs differs from that of Budai-Balke (2009). In Budai-Balke (2009) a possession is defined as a railway track that is unavailable due to maintenance works. The costs resulting from a track possession are mainly dependent on the possession time. In our case, a year in which a replacement is planned can be compared to a track possession. During such a 'possession', a part of the terminal is unavailable for the operational processes for some time. The financial impact of this unavailability is very hard to determine and depends on several factors, e.g. the exact location of the asset, at what time the replacement is executed and the specific combination of replacements. We will therefore use the number of years in which as least one activity is planned, i.e. the number of clusters, as a measure for possession costs. Our goal is therefore to minimize the penalty costs resulting from early or late execution and the number of clusters in the horizon.

In Chapter 4, we will elaborate on the formulation of our model.

4. Model

In this chapter we want to answer the third sub question: *'How can ASM optimally plan the replacements of the assets in the E-pier?'*. As indicated in Chapter 3, our model is based on Budai-Balke (2009) but is modified in order to fit the situation at ASM. In Chapter 4.1, the description of the model is given. In Chapter 4.2, the mathematical formulation of the model is provided.

4.1. Model description

We want to develop a model that balances the costs and benefits of clustering replacements compared to the individual execution of replacements. Here, the direct cost of clustering is the price that is paid for replacing an asset earlier or later than is optimal for this asset. The benefit of clustering is a lower number of clusters. In order to properly balance the penalty costs and the number of clusters in which replacements are planned, a balancing parameter will be used. This balancing parameter represents the relative weight of penalty costs or the number of clusters. Changing this balancing parameter enables the decision maker to steer the model based on the preferences of users or customers in a certain area of the terminal.

In order for our model to decide in which year to schedule activities, several parameters should be known. First, the year the asset was built or the last replacement year before the start of the planning horizon. Furthermore, we have to know the expected lifetime or economic lifecycle of the asset. Based on the construction year or last replacement of the asset and its lifecycle, we can now determine the first replacement moment of the individual asset in the horizon. Likewise, we can determine the subsequent replacement moments. Shifting these individual moments comes at a certain cost, which can be defined for the individual assets. Moreover, we want to be able to bound the allowed deviation from the optimal replacement moment. In specifying these penalty costs and the maximum allowed deviations we can differentiate between early and postponed replacement.

For now ASM decided not to take into account any workload constraints, i.e. it is assumed that the man-hours available for executing the maintenance works are infinite. This is because the work is outsourced and represents external capacity. In the future, the model can however easily be extended to take into account a maximum workload. Moreover, we assume that all replacements can be executed simultaneously. The scheduling of activities within that year in such a way that contractors can manage the workload is outside the scope of this research.

4.2. Mathematical formulation

Assume that we have a planning horizon of $|T|$ and let T be a set of years in which the replacements of an asset a need to be scheduled.

Sets

A	set of assets a
T	years $t \in \{1 \dots T\}$

Parameters

PH	planning horizon
B	balancing parameter that balances the importance of clustering as opposed to the penalty costs
LC_a	the lifecycle of asset a

AE_a the allowable number of years that the replacement of asset a may be executed early

AL_a the allowable number of years that the replacement of asset a may be executed late

Min_a $LC_a - AE_a$, the minimum interval between two replacements of asset a

Max_a $LC_a + AL_a$, the maximum interval between two replacements of asset a

F_a the original first execution moment of the replacement of asset a in the time horizon, without shifting

CE_a the penalty cost for replacing asset a one year earlier than at its end-of-life (linear in the number of years the replacement is planned early)

CL_a the penalty cost for replacing asset a one year later than at its end-of-life (linear in the number of years the replacement is planned late)

Decision variables

$x_{a,t}$ binary variable that denotes if a replacement of asset a is planned in year t (1) or not (0)

y_t binary variable that denotes if at least one replacement is planned in year t ($y_t = 1$ if $\sum_a x_{a,t} \geq 1$; 0 if not)

$EF_{a,t}$ binary variable that denotes if the first replacement of asset a is executed early in year t

$LF_{a,t}$ binary variable that denotes if the first replacement of asset a is executed late in year t

$E_{a,t}$ binary variable that denotes if a replacement of asset a at time t is followed by an early replacement (1) or not (0)

$L_{a,t}$ binary variable that denotes if a replacement of asset a at time t is followed by a late replacement (1) or not (0)

$$Min\left(B\sum_t y_t + \sum_a \sum_t^{t=F_a-1} (EF_{a,t} * CE_a * (F_a - t)) + \sum_a \sum_{t=F_a+1}^{t=F_a+AL_a} (LF_{a,t} * CL_a * (t - F_a)) + \sum_a \sum_t (E_{a,t} * CE_a + L_{a,t} * CL_a) \right)$$

s.t.

$$\sum_a x_{a,t} \leq My_t \qquad \forall t \qquad (1)$$

$$\sum_{t=F_a-AE_a}^{F_a+AL_a} x_{a,t} = 1 \qquad \forall a \qquad (2)$$

$$x_{a,t} \leq M * EF_{a,t} \qquad \forall a, t < F_a \qquad (3)$$

$$
\begin{array}{lll}
x_{a,t} \leq M * LF_{a,t} & \forall a, F_a < t \leq F_a + AL_a & (4) \\
x_{a,t} + \cdots + x_{a,t+Min_a-1} \leq 1 & \forall a, \forall t & (5) \\
x_{a,t} + \cdots + x_{a,t+Max_a-1} \geq 1 & \forall a, \forall t & (6) \\
x_{a,t} + \cdots + x_{a,t+LC_a-1} \leq 1 + M * E_{a,t} & \forall a, \forall t & (7) \\
x_{a,t} + \cdots + x_{a,t+LC_a-1} \geq 1 - M * L_{a,t} & \forall a, t < PH - Max_a + 2 & (8)
\end{array}
$$

The objective function consists of four parts. As explained, the model aims to find the optimum in the trade-off between the penalty costs made for clustering and the number of clusters in which replacements are planned. The first part minimizes the number of clusters. Constraint 1 is related to this first part of the objective function and ensures that the binary variable y_t, which represents the number of clusters, takes the value 1 when there is at least one replacement planned in year t.

The second, third and fourth part of the objective function minimize the costs associated to an early or late replacement of an asset. It can be seen, that these costs are determined in different ways. The reason for this is that the penalty costs associated to the first replacement depend on the deviation from the initial first replacement moment F_a that is given as input to the model, whereas the penalty costs for the subsequent replacements depend on the previous replacement moments, which are not known beforehand. Let us explain how the model works by an example.

The first replacement

Consider $asset1$ with the characteristics as displayed in Table 1.

Lifecycle LC_{asset1}	5
Allowed early AE_{asset1}	1
Allowed late AL_{asset1}	2
Minimum interval Min_{asset1}	5 – 1 = 4
Maximum interval Max_{asset1}	5 + 2 = 7
First moment F_{asset1}	2
Cost early CE_{asset1}	20
Cost late CL_{asset1}	30

Table 1. Characteristics of the fictional asset 1 in our example.

As mentioned before, the initial first replacement of asset a is given by the input parameter F_a. Constraint 2 ensures that the first replacement of asset a is planned somewhere in the allowed interval, based on the value for the initial first replacement of the asset and the allowed early and late execution of the replacement.

The second part of the objective function minimizes the costs associated to an early execution of the first replacement of an asset a. Since we know beforehand the first initial replacement of an asset, we can check rather easily if this replacement is early or late and what the corresponding penalty costs are.

Assume that solving the model has resulted in the planning as in Table 2. The first replacement has apparently been scheduled in year $t = 1$, meaning that the model should take into account $1 *$ €20 in penalty costs.

t	1	2	3	4	…
$x_{asset1,t}$	1	0	0	0	…

Table 2. The first part of the fictional output planning for asset 1 after solving the model.

In constraint 3 it is checked for every $t < F_a$ whether a replacement of asset a is scheduled in year t. Note that for this example, this constraint is only generated for $t = 1$. If a replacement has indeed been scheduled in a year t, the binary early indicator for the first replacement $EF_{a,t}$ is assigned a value 1. If not, $EF_{a,t}$ becomes 0. In the second part of the objective function, the penalty costs for the early first replacements of all assets a are determined. For every t in the interval $[t, F_a - 1]$ the early indicator for the first replacement $EF_{a,t}$ is multiplied by the costs for one year of early replacement CE_a and the number of years that the asset was replaced early, i.e. $F_a - t$. Thus, if the value for $EF_{a,t}$ is 0, the penalty costs generated for this t are also € 0.

Looking at Table 3, we can conclude that the correct penalty costs for early replacement have been generated.

t	$EF_{a,t}$	$F_a - t$	$EF_{a,t} * CE_a * (F_a - t)$
1	1	2 – 1 = 1	1 * €20 * 1 = €20

Table 3. The functioning of constraint 3 and the second part of the objective function.

Now let us look at the penalty costs for a late first replacement. In constraint 4, for every $F_a < t \leq F_a + AL_a$ it is checked if a replacement of asset a is planned in year t. If so, the binary late indicator for the first replacement $LF_{a,t}$ becomes 1. If not, $LF_{a,t}$ becomes 0. Since the first replacement was planned early and not late, no penalty costs for a late replacement are incurred.

t	$LF_{a,t}$	$t - F_a$	$LF_{a,t} * CL_a * (t - F_a)$
3	0	3 – 2 = 1	0 * €30 * 1 = €0
4	0	4 – 2 = 2	0 * €30 * 2 = €0

Table 4. The functioning of constraint 4 and the second part of the objective function.

Looking at Table 3 and 4, it can be seen that indeed the right penalty costs are generated for this first replacement, i.e. €20 for early replacement and €0 for late replacement.

Subsequent replacements

Constraints 5 and 6 plan the subsequent replacements of asset a. Constraint 5 ensures that in every interval $[t, t + Min_a - 1]$ at most 1 replacement of asset a is planned, meaning that it is not allowed to plan the replacement of an asset a at a time t earlier than the last replacement plus the minimum interval. Likewise, constraint 6 ensures that in every interval $[t, t + Max_a - 1]$ at least 1 replacement of asset a is planned, meaning that it is not allowed to plan the replacement of an asset a later than the last replacement plus the maximum interval.

The calculation of the penalty costs of the subsequent replacements is handled in the fourth part of the objective function. Calculating the penalty costs for the subsequent replacements of the asset is somewhat harder than for the first replacement. This is because – for the subsequent replacements – whether the replacement of an asset has been planned early or late depends on when the previous replacement has been planned and we do not know this beforehand.

We can solve this problem by counting how many replacements has been planned in certain intervals. Let us again explain this by an example. Assume that solving the model for a planning horizon $|T| = 15$ results in the planning in Table 5.

t	1	2	3	4	5	6	7	8	9	10	...
$x_{asset1,t}$	1	0	0	0	0	1	0	0	0	0	...

...	11	12	13	14	15
...	0	0	1	0	0

Table 5. The fictional output planning for asset 1 after solving the model.

As we have seen, the first replacement of the asset *asset1* is planned at $t = 1$ and the penalty cost related to this first replacement is already handled by together constraints 3 and 4 and the second part of the objective function.

The second replacement is planned exactly at the asset's end-of-life, i.e. at $t = 6$. The third replacement however planned at $t = 13$ and therefore two years late. The penalty costs for the subsequent replacements of the *asset 1* should therefore be 2 $years\ late * CL_{asset1}$.

In the first column of Table 6 the intervals as generated by constraints 7 and 8 are listed. In the second column we calculated the sum of the values for $x_{a,t}$ over these intervals. In the third and fourth column the values that the early-indicator $E_{a,t}$ and the late-indicator $L_{a,t}$ take for this sum are shown.

Please note that constraint 8 is only generated for $t < PH - Max_a + 2$. If constraint 8 is generated for all values $t \in T$, the late indicator $L_{a,t}$ would wrongly become 1 for intervals generated after the last replacement.

It can be seen that constraint 7 and 8 indeed result in zero years of early replacement and two years of late replacement.

Interval $[\mathbf{t}, \mathbf{t} + \mathbf{LC_a} - \mathbf{1}]$	$\sum_t x_{a,t}$	$E_{a,t}$	$L_{a,t}$	$E_{a,t} * CE_a$	$L_{a,t} * CL_a$
[1, 5]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[2, 6]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[3, 7]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[4, 8]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[5, 9]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[6, 10]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[7, 11]	0	0	1	$0 * €20 = €0$	$1 * €30 = €30$
[8, 12]	0	0	1	$0 * €20 = €0$	$1 * €30 = €30$
[9, 13]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[10, 14]	1	0	0	$0 * €20 = €0$	$0 * €30 = €0$
[11, 15]	1	0		$0 * €20 = €0$	
				$\sum_t (E_{a,t} * CE_a) = €0$	$\sum_t (L_{a,t} * CL_a) = €60$

Table 6. The intervals as generated by constraints 7 and 8, the values for the sum over the values for $x_{a,t}$ over these intervals, and the resulting values for $E_{a,t}$, $L_{a,t}$, $E_{a,t}*CE_a$ and $L_{a,t}*CL_a$.

The model has been programmed in AIMMS. Please see Appendix B for the AIMMS syntax.

4.3. Assumptions and conditions to ensure validity of the model

To ensure validity of the model certain conditions need to be met. In constraints 1, 3 and 4 a Big-M parameter is used. This is a parameter with a large number, compared to the associated variables. Since in constraint 1 the values of $x_{a,t}$ are summed over the entire time horizon, it is safe to take a value of Big-M larger than the number of assets. Furthermore, the costs for an early or late replacement should always be ≥ 1 to ensure a correct functioning of constraints 3, 4, 7 and 8. If the costs for early or late replacement are equal to zero, the early and late indicators may arbitrary become 0 or 1. This may result in a replacement being wrongly indicated as being early or late. Moreover, the year for the first replacement should be smaller than the asset's lifecycle to ensure a correct functioning of constraints 5 and 6.

For constraint 4 to work correctly, the minimum interval Min_a should be larger than the sum of allowed early and allowed late, i.e. $AE_a + AL_a$. If not, the late indicator for the first replacement $LF_{a,t}$ can wrongfully be assigned a value 1 when the second replacement of the asset is scheduled in the interval $[F_a + 1, F_a + AL_a]$. Since the minimum interval Min_a is defined as the lifecycle of an asset, minus the value for allowed early, $LC_a - AE_a$ should be larger than $AE_a + AL_a$. Or: $LC_a > 2AE_a + AL_a$.

Please note that in this formulation of the model, a replacement just before the end of the horizon is not planned whenever it is not necessary, i.e. when the previous replacement + the maximum interval is outside the planning horizon. In other words, a replacement just before the end of the horizon is only planned when the previous replacement + the maximum interval is inside the horizon. The problem is modelled this way since it is assumed that the pier, and therefore the assets within it, loses its value at the end of the horizon. This means it is beneficial to let the assets degrade towards the end of the horizon as much as possible, since this allows ASM to renovate and upgrade the area as a whole without having to devaluate new and healthy assets. In practice, ASM may choose not to perform a replacement just before the end of the horizon in the knowledge that a major upgrade of the area is planned. This choice may then result in postponing the replacement longer than is usually desired (and longer than is dictated by AL_a), more maintenance and therefore higher costs, or preparing for this moment by replacing the asset by a (more expensive) asset with a longer lifecycle. All these options are not taken into account in this research.

5. Case study

This section elaborates on the results of the model and aims to answer the sub question *'What are the benefits of the new methodology?'*. First, in Section 5.1 it is explained how the values for the input parameters were determined. Section 5.2 describes the experimental set up that gives us insight in the performance of the model for different input parameters and discusses the results.

5.1. Input

This section describes how the input parameters for the E-pier case were determined. For every asset we need several input parameters for the model to work with. These parameters are the first replacement F_a, the expected life cycle LC_a of the asset, the allowed number of years of early replacement AE_a, the allowed number of years of late replacement AL_a and the costs per year of early and late replacement CE_a and CL_a. We will shortly describe how the values for these parameters have been determined.

Aggregation level

The database with E-pier assets consist of more than 2,000 assets. Gathering reliable data for every individual asset might prove itself to be very hard and time-intensive task. For this research, the level of detail was therefore chosen to be on asset type. This means that the values for the input parameters LC_a, AE_a, AL_a, CE_a and CL_a are equal for all individual assets of the same asset type and are based on averages. This allows us to aggregate the data on asset type and construction year. The E-pier for example has 17 elevators, two of them built in 1999, two in 2003 and 13 elevators built in 2015. Thus, instead of 17 individual elevators, we can aggregate them in three asset groups. This aggregation of individual assets on asset type and construction year results in only 116 asset groups to be planned, which also significantly speeds up the computation time when solving the model. Where solving the model with aggregated data takes only 8.136 seconds, solving the model for the individual assets takes 424.66 seconds.

Please note that the term 'construction year' might also represent the latest year the asset was replaced. Based on these construction years the first replacement year F_a is determined by adding the expected lifecycle to the construction year, i.e. $construction\ year + LC_a$. In case $F_a < 2019$, F_a is set to 2019.

Penalty costs (CE_a and CL_a)

The penalty costs for one year of early or late replacement are determined based on the replacement costs of the asset. This replacement cost is based on averages over the individual assets of that asset type. For example, the costs for the replacement of an elevator are estimated to be €185,000 per elevator. These are the costs that are budgeted by the asset owner for the replacement of a standard elevator with one cabin access point, three shaft accesses and three stopping points. In reality the replacement cost may differ. When the elevator has for example one additional stopping point, the replacement cost will be higher. For more detailed information on how the replacement costs for the different asset types are determined and the assumptions that have been made, please see Appendix C.

The replacement cost for each group of assets depends on the number of assets in that specific group. The groups 'elevators 1999' and 'elevators 2003' for example both consist of 2 assets, resulting in replacement costs of $2 * €185{,}000 = €370{,}000$. The group 'elevators 2015' consists of 13 assets, resulting in replacement costs of $13 * €185{,}000 = €2{,}405{,}000$. The costs for an early and late replacement are based on the fact that assets are linearly depreciated over their expected lifetime. The costs for one year of early replacement CE_a are therefore set equal to the total replacement costs divided by the lifecycle, i.e. $CE_a = \frac{total\ replacement\ costs}{LC_a}$.

The costs for a late replacement CL_a are hard to determine. ASM however feels that replacing an asset late is more expensive than replacing an asset early, since maintenance costs, obsolescence of spare parts and risks on failure of the asset tend to increase when replacement is postponed. Therefore we asked the Technical Management team to estimate how much more expensive or undesired one year of late replacement would be relative to one year of early replacement. Based on expert judgements we decided to differentiate between 10% and 20% additional costs. We decided to work with 20% additional costs for asset types for which the costs and risks are believed to increase significantly over time. This are the mechanical assets with moving parts, like elevators, escalators and closing installations. For assets that show less degradation over time, for example building elements, we decided to work with an additional 10% as compared to the costs for one year of early replacement. For the replacement values and penalty costs for all asset types, please see Appendix A.

Thus, getting back to our elevator example, for the group 'elevators 1999' with a replacement cost of €370,000, one year or early replacement would cost ASM $CE_a = \frac{€370{,}000}{25} = €14{,}800$. One year of late replacement CL_a would then cost ASM $1{,}2 * €14{,}800 = €17{,}760$. Like the costs for an early replacement (because of the linear depreciation that ASM uses), the costs for late replacement are assumed to be linear in the number of years that an asset is replaced late. Replacing the group 'elevators 1999' two years early would therefore cost ASM $2 \times CE_a = 2 \times €14{,}800 = €29{,}600$. Replacing this asset group two years late costs $2 \times CL_a = 2 \times €17{,}760 = €35{,}600$.

Allowed early- and late replacement (AE_a and AL_a)

Also for determining the allowed number of years an asset might be replaced early or late, we consulted with Technical Management. It was decided to differentiate between critical and less critical assets, where critical assets may only be postponed with 10% of their lifecycle, whereas less critical assets may be postponed with 20%. To assess the criticality of an asset type, several factors are taken into account. First, the extent to which an asset – or failure of an asset – directly influences to the airport's primary processes. Also whether or not an asset fulfills a safety function contributes to its criticality. Moreover, an asset's importance to Schiphol's strategic goals is taken into account. Important topics here are sustainability and an excellent visit value. For assessing the criticality of the assets in the E-pier database, the document '*Beslismodel kritieke assets*' was used. In this document the factors mentioned above are evaluated for all of AAS's assets.

Examples of assets important to the primary process are elevators, escalators, closing installations and low voltage installations. An example of an assets with a safety function are back-up power units. Asset types important in Schiphol's strive for sustainability goals are for example air treatment units, using relatively much energy. Important to an excellent visit value are for example moving walkways and climate systems. Passengers do not like to walk long distances and should not get cold when sitting down for a couple of hours, waiting for their airplane to depart. Also important are certain building- and architectural elements. ASM believes that the decay and ageing of non-constructive and finishing building elements – as for example floor tiles and paintwork – is harmful to the customer experience.

Less critical assets, i.e. assets that do not directly influence the primary processes at AAS, do not have a safety function and do not directly contribute to Schiphol's strategic goals may be postponed with 20% of their lifecycle. Examples of these assets are constructive building elements and rain- and wastewater drainage installations.

Technically, the allowed number of years do not necessarily have to be bounded, since the disadvantage of an early replacement is represented by the costs for early depreciation and is therefore already covered by the penalty costs for an early replacement CE_a. Bounding the number of years however speeds up the model, since it limits the number of combinations. Therefore, for now we set the allowed years that an asset may be replaced early to be equal to $1{,}5 * AL_a$, i.e. an asset may always be replaced 50% earlier than it may be postponed. In Section 5.5 we will come back to this to see how this decision influences the model. The allowed years of early and late replacement for all asset types, are also included in Appendix A.

5.2. Results for different values of the balancing parameter

The value for the balancing parameter is of significant importance to the model: the higher the value, the more weight on minimizing the number of clusters and the more penalty costs are made to achieve this. We have experimented with the value of the balancing parameter to see what output the model generates for different values. This section also aims to give a deeper insight in the consequences of clustering and provides information on the characteristics of the final planning of which the Technical Management team thinks they are important for final decision making. This should eventually help ASM in making a deliberate choice in how many penalty costs to accept.

The goal of this research is to find out how ASM can better cluster its replacements over the long-term. We have seen that in the baseline situation, i.e. the scenario in which replacements are always executed at the end of an asset's lifecycle, replacements are planned in 51 of the 60 years. The number of years with which we can decrease the number of clusters depends on the height of balancing parameter. The experiment was run with the value for the balancing parameter varying from 1,000 to 200,000 with a step size of 1,000. Each configuration results in a planning with a corresponding value for the number of clusters and penalty costs. When calculating the relative savings, the resulting number of clusters is compared to the number of clusters in the baseline situation (51). Thus, decreasing the number of clusters to 49 equals relative savings in number of clusters of $\frac{51-49}{51} = 3{,}92\%$. Table 7 summarizes the output of the model for the different values for the balancing parameter.

INPUT	**OUTPUT**		
Balancing parameter $\times \mathbf{10^3}$	**Penalty costs**	**Number of clusters**	**Relative savings in number of clusters**
1 – 2	€0	51	0%
3 – 4	€4,706	49	3.92%
5	€8,765	48	5.88%
6	€34,095	43	15.69%
7	€59,707	39	23.53%
8 – 9	€74,867	37	27.45%
10 – 11	€92,982	35	31.37%
12 – 13	€116,316	33	35.90%
14 – 36	€226,802	25	50.98%
37 – 47	€441,651	20	60.78%
48 – 67	€552,907	17	66.67%
68 – 108	€687,443	15	70.59%
109 – 145	€904,043	13	74.51%
146	€1,049,135	12	76.47%

Table 7. The output of the different values for the balancing parameter.

As can be seen in Figure 5, there seems to be a strong logarithmic relationship between the penalty costs the decision maker is willing to accept and the relative savings in number of clusters achieved. In the left part of the graph, the relative savings rapidly increase as the penalty costs the decision maker wants to accept increase. As the accepted penalty costs increase, the savings keep increasing but more slowly. In other words, the additional costs the decision maker would have to accept for one cluster less exponentially grows. The improvement from 51 to 49 clusters would cost ASM €4,706, i.e. €2,353 per year, whereas the improvement from 13 to 12 clusters one should accept an additional €145,092 in penalty costs.

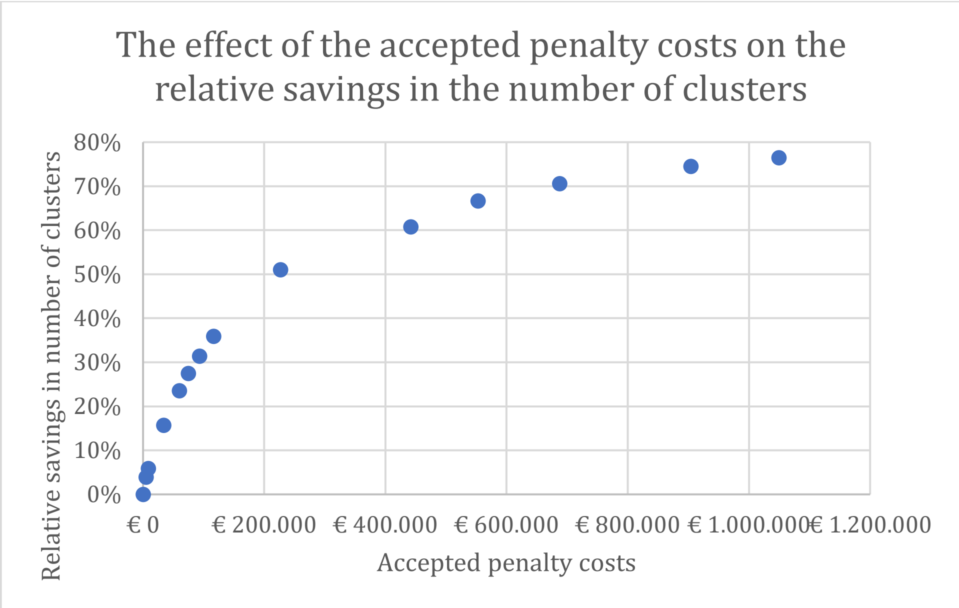

Figure 5. The relationship between the penalty costs the decision maker chooses to accept and the relative savings in the number of clusters.

For the input of the E-pier case as in Appendix A, the minimum number of clusters that one can achieve is 11 clusters over a horizon of 60 years. This was tested by setting the penalty costs for an early or late replacement to its minimum, i.e. €1, and assigning a very high value to the balancing parameter (5×10^9). The maximum relative savings for the E-pier case therefore equal $\frac{(51-11)}{51} * 100 = 78{,}43\%$. When ASM's goal is to minimize the number of clusters at all costs, a decrease from 51 to 11 clusters can be achieved by making €4,926,934 in penalty costs for early and late replacements.

Which solution is optimal depends on the penalty costs ASM wants to accept, but other metrics are possibly even more important the costs. Investing €1,049,135 in penalty costs in order to decrease the number of clusters with 76.47% is acceptable to ASM, but this depends on the specific decisions that are made by the model. Maximally postponing all replacements might for example not represent an acceptable solution. Therefore other important indicators of the quality of the resulting planning are the percentage of assets that have been planned early or late, the number of years with which assets have been planned early or late and the distribution over replacements values per cluster.

In Figure 6 we see the relationship between the penalty costs the decision maker chooses to accept and the number of years asset groups were replaced early or late. The left y-axis represents the percentage of assets that was replaced early or late. In every planning, if an asset was replaced early or late this only happened once over the horizon. A percentage of 50% early replacement therefore means that 50% of the asset groups (i.e. $0.5 \times 116 = 58$ asset groups) were replaced early once. The y-axis at the right represents the average number of years asset groups were replaced early or late. This average only included the assets that were indeed shifted. So, for example in the situation in which ASM accepts €1,049,135 in penalty costs, asset groups that were replaced early were on average shifted with 1.10 years, whereas asset groups that were replaced late were on average shifted with 1.29 years.

Figure 6. The relationship between the penalty costs the decision maker chooses to accept and the number of years or early/late replacement: at the left y-axis the percentage of asset groups that is replaced at least one time early or late over the planning horizon, at the right y-axis the average number of years with which assets were shifted.

As expected, the percentage of asset groups that are shifted increases when the accepted penalty costs increase. It can be seen that the percentage of assets that is replaced late stays relatively stable, whereas the percentage of assets that is replaced early rapidly grows. This can directly be explained by the way we defined our cost function, i.e. the fact that a late replacement was always more expensive than an early replacement. It is however the exact opposite of what ASM does now and we can therefore ask ourselves if this indeed properly represents the actual situation. We will come back to this in Section 5.4.5 where we propose a non-linear cost function.

In Figure 6 we saw the average number of years with which assets were replaced. Let us further examine this. In Figure 7 the number of years with which an asset group is replaced late or early is displayed. When the accepted penalty costs become €116,316 or more, we see that some asset groups are replaced two years early. This is however very limited. Most asset groups are only replaced only one year early. From €226,802 in penalty costs onward, some asset groups are replaced two years late. It can however be concluded that relatively few asset groups are replaced late and that a replacement of two years late is exceptional. This is an interesting finding since it shows us that it is not necessary to long postpone replacements in order to decrease the number of clusters and that the possible risks that come with clustering are limited. Therefore, the solution planning as generated for all different values of the balancing parameter is acceptable to ASM. It can be concluded that by making relatively small shifts of only one or sometimes two years, the number of clusters can already be heavily reduced.

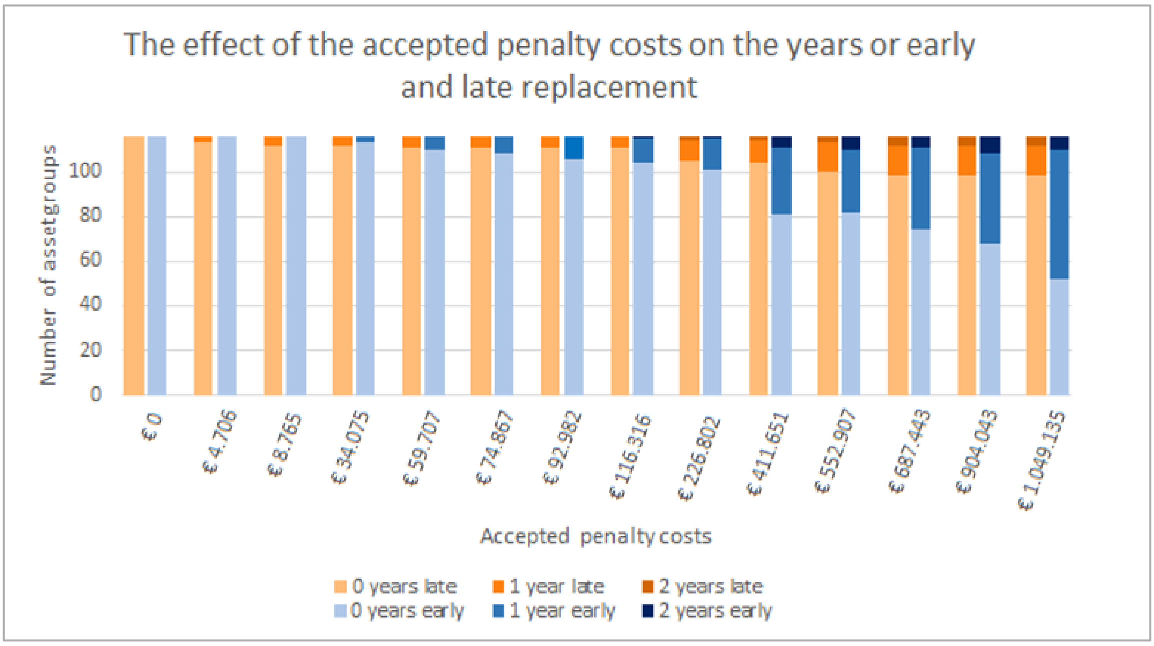

Figure 7. The relationship between the penalty costs the decision maker chooses to accept and the number of asset groups that are replaced 0, 1 or 2 years late (in orange) and early (in blue), represented in a stacked bar chart.

Interesting to see also is that most of the penalty costs flow from the early or late planning of the first replacement, and – in most cases – only a small percentage of the penalty costs comes from the early or late replacement of the subsequent replacements. Apparently it pays off to shift replacement moments early in the horizon, such that later in the horizon replacements are synchronized in a natural cadence. This is logical, since when replacements are aligned later in the horizon, this can easily result in more clusters than necessary. Take for example two assets with a lifecycle of three years. When the first initial replacement of asset A is scheduled in year 2, and the first initial replacement of asset B in year 3, not shifting the assets would result in 7 clusters over a 10 year horizon (see Table 8).

t	1	2	3	4	5	6	7	8	9	10	...
$x_{assetA,t}$	0	1	0	0	1	0	0	1	0	0	...
$x_{assetB,t}$	1	0	0	1	0	0	1	0	0	1	...

Table 8. Fictional baseline situation, i.e. assets are replaced immediately at their end-of-life, for assets A and B.

When we however shift the first replacement of either asset A 1 year earlier, or asset B 1 year later (whichever is cheaper), the replacements of the assets are directly lined-up, resulting in only 4 clusters (see Table 9).

t	1	2	3	4	5	6	7	8	9	10	...
$x_{assetA,t}$	1	0	0	1	0	0	1	0	0	1	...
$x_{assetB,t}$	1	0	0	1	0	0	1	0	0	1	...

Table 9. Fictional clustering for assets A and B, resulting in 4 clusters.

When not the first, but the second replacement was shifted, this would have resulted in 5 clusters (see Table 10). This explains why shifting the first replacement is often beneficial. This is however not always the case. Sometimes the replacements of assets may come closer to each other with the passing of time and it may be beneficial to wait with aligning them until their replacement moments are closer. This is more likely to happen when clustering is less important, i.e. when the balancing parameter is relatively low, since although the associated cost will be lower, it will also result in more clusters.

t	1	2	3	4	5	6	7	8	9	10	...
$x_{assetA,t}$	0	1	0	1	0	0	1	0	0	1	...
$x_{assetB,t}$	1	0	0	1	0	0	1	0	0	1	...

Table 10. Fictional clustering for assets A and B, resulting in 5 clusters.

Figure 8 shows the proportion of the total penalty costs that comes from early or late replacement of the first moment and the subsequent moments.

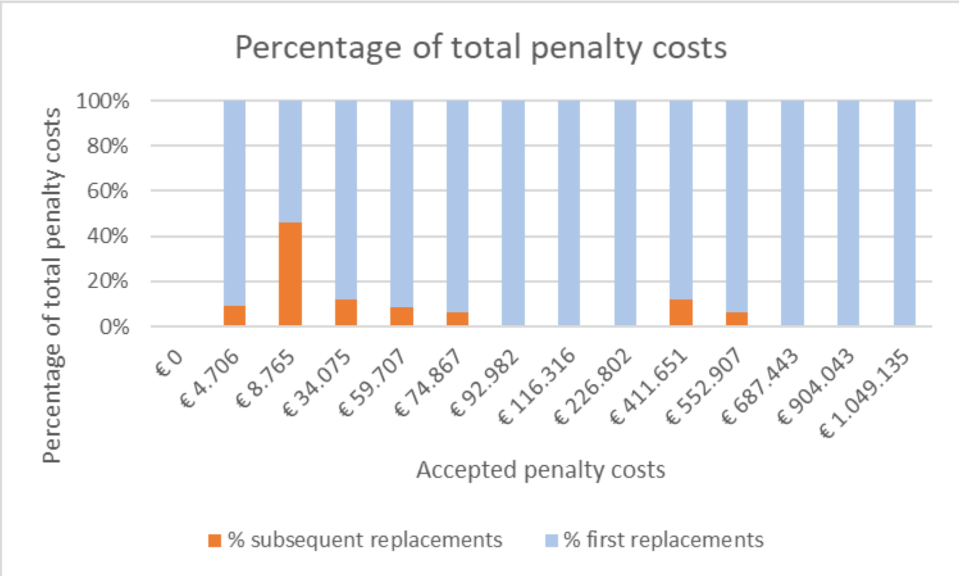

Figure 8. The proportion of the total penalty costs that comes from early or late replacement of the first moment and the subsequent moments.

For the validation of the model and the sensitivity analysis, we will compare different scenarios and analyze them in more detail. Based on the outcomes presented in this section, the Technical Management team has showed interest in seeing the results of three scenarios compared in more detail. We have already seen one of them, i.e. the baseline situation. In this scenario no penalty costs are accepted, which means all assets are replaced at their individual moments, resulting in 51 clusters. The other two scenarios would be Scenario 1 in which ASM accepts €226,802 in penalty costs, resulting relative savings of 50.98% (25 clusters), and Scenario 2 in which ASM accepts €1,049,135 in penalty costs, resulting in relative savings of 76.47% (12 clusters). The different scenarios are summarized in Table 11 and Table 12.

Scenario	Penalty costs	Number of clusters	Relative savings in number of clusters
Baseline situation	€ 0	51	0%
Scenario 1	€ 226,802	25	50.98%
Scenario 2	€ 1,049,135	12	76.47%

Table 11. The different scenarios, with the corresponding penalty costs, the number of clusters and relative savings.

Scenario	Early replacements	Average years early	Late replacements	Average years late
Baseline situation	0	0	0	0
Scenario 1	14	1.07	9	1.22
Scenario 2	58	1.10	14	1.29

Table 12. The different scenarios, with the corresponding number of assets that are replaced early or late and the average number of years with which shifted assets were replaced early or late.

For these scenarios the replacement values per cluster have been displayed in more detail. We have already seen Figure 9 in section 2.3 (Figure 4). This figure shows the distribution of the replacement value over the planning horizon for the base scenario. Figures 10 and 11 show these distributions for scenarios 1 and 2 respectively.

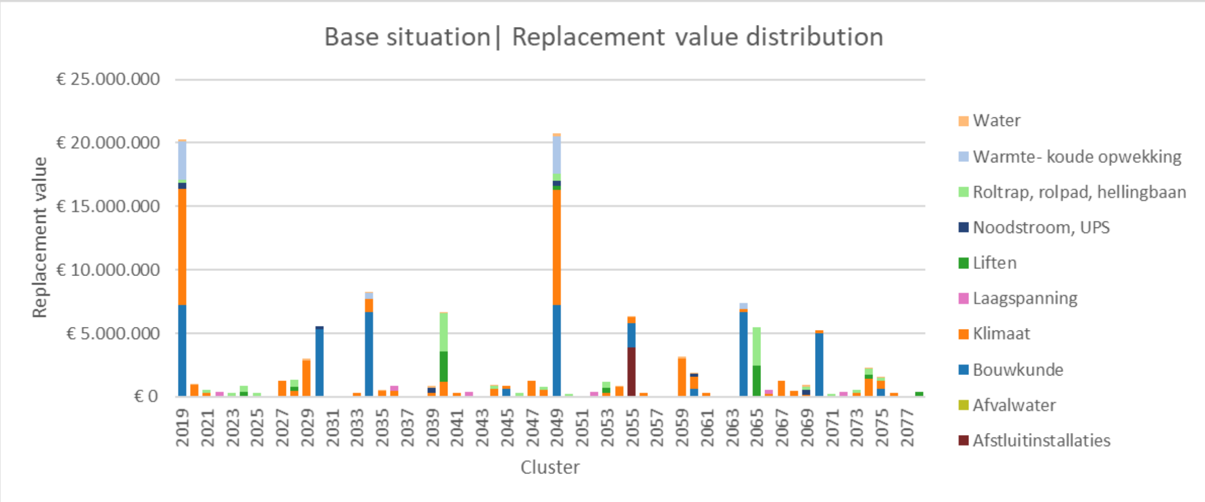

Figure 9. The distribution of the replacement values for the replacements planned in each cluster for the baseline situation.

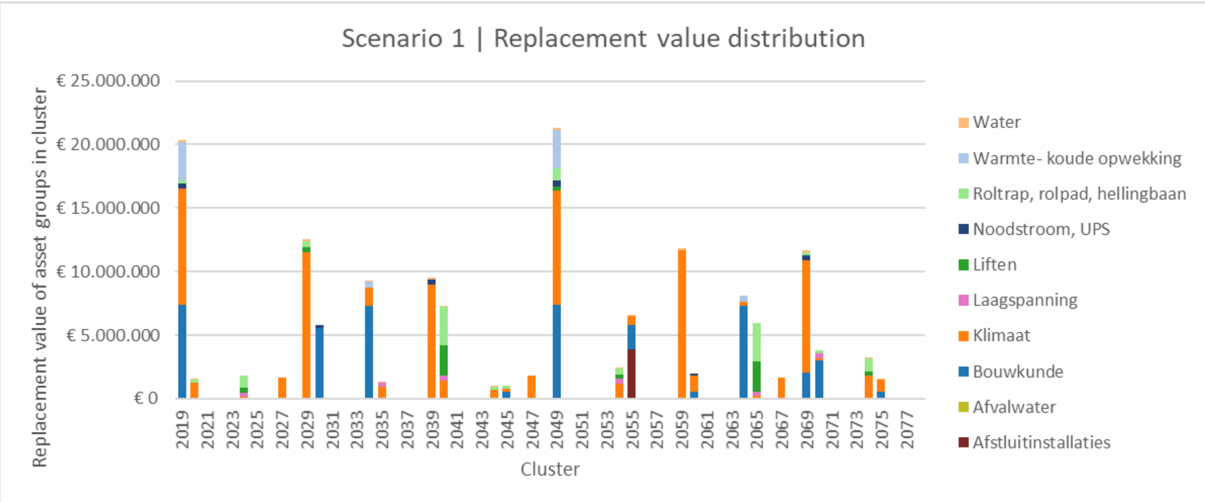

Figure 10. The distribution of the replacement values for the replacements planned in each cluster for Scenario 1.

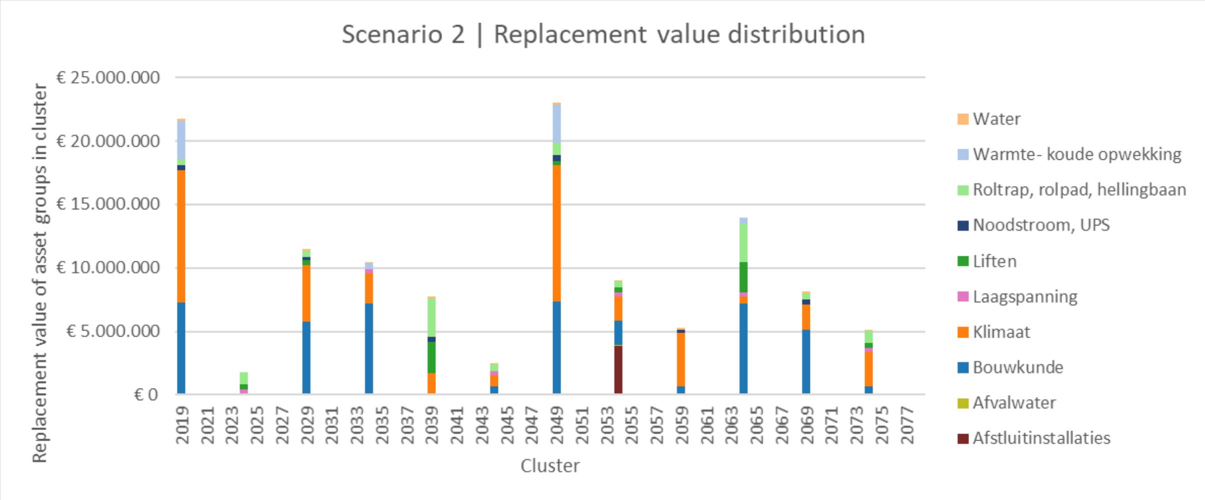

Figure 11. The distribution of the replacement values for the replacements planned in each cluster for Scenario 2.

Now that we know the solution it may seem straightforward, and one may ask what the added value of the model is compared to intuitively clustering the replacements. When we compare the solutions of scenarios 1 and 2 (in Figure 10 and 11) to that of the baseline situation (Figure 9), we can conclude however that the decisions made by the model are sometimes counterintuitive. The baseline situation for example shows peaks in 2040, 2055 and 2070. One may therefore expect to plan clusters in these years. The solution for scenario 2 however clearly shows that it is better to plan the replacements around these years in 2039, 2054 and 2069 respectively. This is apparently beneficial on the long-term and limits the number of clusters. Moreover, one may expect only to shift asset groups with a relatively low replacement value. Although asset groups for which it is cheaper to shift are indeed shifted more often and further, many more factors are important in the decision in which year to replace a group. Very important here is how the choice for a replacement moment interacts with future replacements.

Part of the solution at least as important to ASM as the penalty costs, the number of clusters and the number of assets that have been shifted, is the answer to the question which assets should be replaced in which year, i.e. the planning itself. It should be noted however that these solutions are non-unique, meaning that there may be other solutions that result in the same penalty costs and the same number of clusters.

5.3. Computing time

The formulated problem is a combinatorial one. We can calculate the number of possible combinations by multiplying the number of an asset group's replacements in the horizon ($60/LC_a$) by the number of years in which a replacement is allowed to be planned ($1 + AE_a + AL_a$). The resulting number represents the possible number of years in which a replacement of the asset group may be planned. When we multiply this possible number of years in which a replacement of the asset can be planned for all assets, we find that there are 2.79×10^{146} possible combinations of replacement moments for the input parameters as defined in this research. Also when we would allow each replacement to be planned only one year early and one year late, there are already 2.64×10^{100} possible replacement combinations. For the database used in this case one might be able to manually construct a long-term replacement planning that approaches the objective function value that is found by the model, but when the model becomes more complex (e.g. more variation in first replacement moments F_a and lifecycles LC_a) it becomes harder to find a good solution manually and the model will significantly outperform a pragmatic, intuitive planning.

The model is remarkably fast and able to find the optimal combination of replacements in 2.911 seconds for Scenario 1 and 8.136 seconds for Scenario 2. As we already mentioned in Section 5.1, Scenario 2 runs in 424.66 seconds when not aggregating the data beforehand. It should be noted however that this computing time significantly increases when the problem becomes more complex, for example when there is more variation in the lifecycles of the assets or in the first replacement years. It is important to limit Big-M to the number of assets + 1, since the larger Big-M becomes, the more it unnecessarily slows down the model. Also increasing the balancing parameter increases the computing time.

5.4. Sensitivity analysis

The E-pier at AAS was taken as a case for this research and although we believe the assets included in the dataset well represent an average pier at AAS, we want to know how and to what extent changes in the data influence the solution. In addition, for some input parameters assumptions have been made. In this section we will challenge these assumptions and find out how and to what extent the solution depends on the input parameters. In order to do so, we will run different experiments and see how they influence the solutions in scenarios 1 and 2 in specific. Please recall that scenario 1 and 2 represent the situations in which the number of clusters is decreased to 25 (50.98% savings) and 12 (76.47% savings) respectively.

5.4.1. Allowed early (AE_a) and allowed late (AL_a)

In the original situation the allowed years a replacement may be postponed AL_a was set to either 10% or 20% of an asset's lifecycle and AE_a was set to $1{,}5 \times AL_a$. We want to see how this decision influences the solution planning.

Dismissing the differentiation and setting all values for AL_a to be either 10% (experiment 1a) or 20% (experiment 1b) for all assets does not influence the solution.

Experiment	Description
Original	AL_a = 10% or 20% of an assets' lifecycle (as described in section 5.1) $AE_a = 1{,}5 \times AL_a$
1a	AL_a = 10% for all asset groups
1b	AL_a = 20% for all asset groups
1c	$AL_a = AE_a = 1$ year

Table 13. The different configurations of experiment 1.

As expected, when we limit AE_a and AL_a to only one year the output of the model and the objective function value do change. When ASM now wants to decrease the number of clusters to 25 (Scenario 1), it should be willing to accept penalty costs of €373,749 instead of €226,802. With these values for AE_a and AL_a, the solution of 12 clusters (Scenario 2) cannot be reached anymore. This was tested by assigning a very high value to the balancing parameter (5×10^9), and setting the costs for early and late replacement to their minimum values, i.e. €1. The minimum number of clusters that can be reached now equals 16 clusters. To reach this, ASM would have to accept €3,049,235 in penalty costs. The results of experiment 1c are summarized in Table 14.

Scenario	Experiment	Penalty costs	Penalty costs Δ	Number of clusters
1	Original	€226,802	-	25 (50.98%)
1	1c	€373,749	+ 64.7%	25 (50.98%)
2	Original	€1,049,135		12 (76.47%)
2	1c	€3,049,235	+ 190.6%	16 (68.63%)

Table 14. The results of experiment 1c.

We see that the penalty costs significantly increase when we limit AE_a and AL_a to only one year, i.e. with 64.7% and 190.6% for scenario 1 and 2 respectively. In Section 5.3 we saw that in the original solutions of both scenarios 1 and 2, all shifts in replacements were made when planning the first replacements, such that future replacements rapidly became synchronized in a natural cadence. When we limit the number of early and late replacements to one year, we see that also subsequent replacements need to be shifted. In addition, the replacements of many asset groups have to be shifted multiple times during their lifecycle whereas this was never needed in the original solution. The number of early and late replacements and the total number of years groups of assets are replaced early or late is summarized in Table 15. Columns 3 and 5, 'Early replacements' and 'Late replacements', refer to the number of times that asset groups were planned early or late. Columns 4 and 6, 'Total years of early replacement' and 'Total years of late replacement', give the total number of years with which assets were shifted. In Scenario 2, experiment 1c for example, 93 asset groups were replaced early. The total years of early replacement however equals 111. This indicates that some assets were shifted multiple times over the time horizon, i.e. not only in the first replacement but also in the subsequent ones. Looking at Table 15, we can see that in order to reach the same decrease in the number of clusters, many additional shifts are needed, which results in higher penalty costs.

Scenario	**Experiment**	**Early replace-ments**	**Total years of early replacement**	**Late replace-ments**	**Total years of late replacement**
1	Original	14	15	9	11
1	1c	22	22	33	41
2	Original	58	65	14	18
2	1c	93	111	64	67

Table 15. The number of early and late replacements, as well as the number of years assets are replaced early or late significantly increases when AL_a and AE_a are limited to 1.

This experiment shows us that limiting the allowed number of years of early and late replacement can result in significantly higher penalty costs, since more shifts need to be made in order to achieve the same decrease in the number of clusters. Synchronizing the replacements early in the horizon is therefore beneficial since this way penalty costs can be kept relatively low. ASM is advised to see if it is indeed acceptable for all assets to shift them with 2 years. If not, the solution may change.

5.4.2. Already fully depreciated assets ($F_a < 2019$)

We have seen that the first replacement moment for many assets lies in the past (25% of the assets in this case). For these assets, in the original situation the first replacement year was set to 2019, representing a catch-up. This decision is likely to heavily influence the solution planning and practically impossible. Another way to deal with the backlog is spreading those replacements out over a longer period of time.

Let is now assume that instead of setting the initial first replacement year to 2019, the condition of these fully depreciated assets is inspected and based on this condition the initial replacement moment is determined to be somewhere in the next five years. What would this mean for the solution and the penalty costs? We have simulated this situation by assigning a random first replacement year between 2019 and 2024 to those assets.

Important to note is that, since these assets are already fully depreciated, we do not want to charge any costs in case the first replacement is planned earlier than the initial first replacement. This would namely not represent a disinvestment. Therefore, for these assets EF_a should be zero. For a late first replacement LF_a we do want to charge a penalty. For the subsequent replacements of these assets penalty costs for both early and late replacement are charged, similar to all other assets.

Since we work with random values here, the experiment was repeated ten times for both Scenario 1 and Scenario 2. The resulting solutions differ from the original solution. More assets need to be replaced late in order to reach the same decrease in the number of clusters. As we also saw in Section 5.3.1, also in this experiment many subsequent replacements of assets need to be shifted, instead of only their first replacement moment as was the case in the original situation. The average penalty costs for scenario 1 and 2 are €277,732 and €1,204,315 respectively, as summarized in Table 16.

Scenario	Experiment	Penalty costs	Penalty costs Δ	Number of clusters
1	Original	€ 226,802	-	25 (50.98%)
1	2	€ 277,732	+ 22.5%	25 (50.98%)
2	Original	€ 1,049,135	-	12 (76.47%)
2	2	€ 1,204,315	+14.8%	12 (76.47%)

Table 16. The results for experiment 2.

Interesting to see is that in this experiment, the solution for Scenario 2 over time converges to the original solution but it takes more time to reach a steady cadence than in the original solution. In the original solution of Scenario 2 there is a cadence in which clusters are planned every five years directly from the start of the horizon. In this experiment, the clusters are less evenly distributed in the beginning of the horizon. After some time however, the planning follows the cadence in which replacements are planned every five year, equal to the original planning for this scenario.

5.4.3. Replacement value for constructive assets

The costs for replacing an asset group one year early or late depend on the asset's group replacement costs. One might ask if and how the solution changes when these replacement values are higher or lower than expected. First it should be noted that when the costs are higher or lower on all fronts, the solution will not change. Increasing or decreasing all replacement values with X% will just respectively increase or decrease the total penalty costs with the same X%.

Important to the model are the relative penalty costs for an asset group, i.e. how expensive shifting the replacement of this group is compared to the shifting of other asset groups. All other things being equal, the model will choose to shift the asset group that will generate the lowest penalty costs. Although for most asset types reliable information regarding replacement values was available, for the estimation of the replacement costs of constructive- and building elements many assumptions had to be made. Therefore, we investigated if and how the solution changes when the replacement costs of constructive asset types change.

In experiments 3a – 3d, the replacement costs for these assets have been set to respectively 80%, 90%, 110% or 120% times the costs in the original situation. This does not change the solution, i.e. clusters are being planned in the same year. The only thing that changes are the penalty costs. These are obviously lower (in experiments 3a and 3b) or higher (in experiments 3c and 3d) than in the original solution. For the penalty costs resulting from running these experiments, please see Appendix D.

5.4.4. Balance CE_a and CL_a

The penalty cost for one year of early replacement CE_a was calculated as the replacement value of an asset group divided by the depreciation period. The penalty cost for one year of late replacement CL_a was then determined by multiplying CE_a by a factor 1.1 or 1.2 based on expert judgement (see Section 5.1). We are interested in what happens when we change these factors.

Not differentiating between a factor 1.1 and 1.2, but instead multiplying all values for CE_a by the same factor may theoretically change the original solution, since it changes the relative costs of late replacement of the assets towards one another. In our case, it did not affect the original solution. So when does the solution change? In the original situation replacing an asset three years early is always more expensive than replacing an asset two years late. The model will therefore, for this dataset, never replace an asset three years early. This starts to change when the costs for one year of late replacement become more than two times as big as the costs for one year of early replacement, i.e. when $CL_a > 2 \times CE_a$. In this case, the model will make different decisions in order to decrease the number of clusters to 25 (in Scenario 1) or 12 (in Scenario 2). Table 17 shows the resulting penalty costs. Table 18 shows that, as expected, the number of late replacements decreases and the number of early replacements increases. Columns 3 and 5 again present the number of assets that were replaced early or late. In column 4 and 6 it is displayed with how many years assets were shifted on average, given they were shifted.

Scenario	Experiment	Penalty costs	Penalty costs Δ	Number of clusters
1	Original	€ 226,802	-	25 (50.98%)
1	4	€ 297,907	+ 31.6%	25 (50.98%)
2	Original	€ 1,049,135	-	12 (76.47%)
2	4	€ 1,203,860	+14.7%	12 (76.47%)

Table 17. The resulting penalty costs for experiment 4.

Scenario	Experiment	Early replacements	Average years early	Late replacements	Average years late
1	Original	14	1.07	9	1.22
1	4	16	1.31	7	1
2	Original	58	1.12	14	1.29
2	4	62	1.24	10	1

Table 18. The results for experiment 4 in terms of early and late replacements.

5.4.5. Non-linear penalty costs for late replacements

In Section 5.2 we saw that the vast proportion of shifts in replacements came from early replacement and that relatively little assets were replaced late. This can easily be explained since one year of late replacement was always more expensive than one year of early replacement. This is however the opposite of what ASM does now. In this section we therefore want to propose another cost structure for the situation in which assets are replaced late that may better represent reality.

In the original model we assumed the penalty costs for a late replacement to be linear in the number of years with which assets are replaced late. In reality, these costs are likely to increase over time and therefore to be non-linear. Replacing an asset only one year late may then be cheaper than one year of early replacement, but these costs will increase with the number of years with which the replacement is postponed and at some point replacing an asset late will become more expensive than early replacement.

Modelling these non-linear penalty costs for late replacement requires us to make some modifications in the original model as it was formulated in Section 4.2. The most important change is found in constraint 8. Please recall that in the original model constraint 8 checks for all intervals of length $[t, t + LC_a - 1]$ how many replacements are scheduled. If no replacements are scheduled in the interval, the binary late indicator $L_{a,t}$ gets assigned a value 1. This constraint therefore counts how often a replacement is replaced at least one year late and, as we showed in Table 6, generated the correct penalty costs. This approach was sufficient in case of linear increasing penalty costs, but using this approach we cannot differentiate between the number of years with which an asset was replaced late. We can therefore not penalize a two-year late replacement heavier than we would penalize a one-year late replacement, as would be the case with non-linear costs. We therefore have to introduce an additional indicator that denotes the number of years with which an asset was replaced late. We can do this by replacing constraint 8 by the following constraint:

$$x_{a,t} + \cdots + x_{a,t+LC_a-1+(u-1)} \geq 1 - M * L_{a,t,u}$$

The binary late indicator $L_{a,t,u}$ now becomes 1 when a replacement was replaced at least u years late. We can now determine the costs for late replacement of an asset as follows:

$$CL_a * L_{a,t,u} * u$$

By defining the costs like this, the costs are quadratic in the number of years with which an asset is replaced late. Let us see why. Suppose that after running the model the generated planning for the replacements of an asset is as shown in Table 19. Let us say this asset has a lifecycle LC_a of two years, which means the second replacement was planned two years late. In Table 20 the intervals as generated by this new constraint are displayed, as well as the corresponding values for the new late indicator $L_{a,t,u}$.

t	1	2	3	4	5	...
$x_{a,t}$	1	0	0	0	1	...

Table 19. The first part of the fictional output planning for an asset after solving the model.

u	Interval $[t, t + LC_a - 1 + (u - 1)]$	$\sum_t x_{a,t}$	$L_{a,t,u}$	$CL_a * L_{a,t,u} * u$
1	[1, 2]	1	0	€10 * 0 * 1 = €0
1	[2, 3]	0	1	€10 * 1 * 1 = €10
1	[3, 4]	0	1	€10 * 1 * 1 = €10
1	[4, 5]	1	0	€10 * 0 * 1 = €0
2	[1, 3]	1	0	€10 * 0 * 2 = €0
2	[2,4]	0	1	€10 * 1 * 2 = €20
2	[3,5]	1	0	€10 * 0 * 2 = €0
				$\sum_u \sum_t (CL_a * L_{a,t,u} * u) = €40$

*Table 20. The intervals as generated by the new constraint, the values for the sum over the values $x_{a,t}$ over these intervals, and the resulting values for $L_{a,t,u}$ and CL_a*$L_{a,t,u,}$*u.*

When the asset was replaced only one year late, i.e. at $t = 4$, only $L_{a,2,1}$ would be 1, therefore resulting in €10 in penalty costs. When the replacement was three years late, i.e. at $t = 6$, three more late indicators, i.e. $L_{a,4,1}$, $L_{a,2,2}$ and $L_{a,2,3}$ would have become 1, resulting in €90 in penalty costs. Thus, modelling the problem this way, the cost for a late replacement of u years equals $CL_a * u^2$. For the modified model formulation in case of non-linear costs, please see Appendix E.

Please recall that in Section 5.1 we explained that the costs for late replacement CL_a were set equal to the costs for a year of early replacement CE_a, but by adding either 10% or 20% in order to account for additional maintenance costs, obsolescence of spare parts and risks of failure. We will now apply the same differentiation, but in a different way. We will now set CL_a to either 1% or 2% of an asset group's replacement value. By doing this, for all assets one year postponing becomes cheaper than one year of early replacement. For 57 of the 116 asset groups (49.14%), two years postponement is still cheaper than one year of early replacement. For 11 of the 116

(9.5%) of the asset groups three years postponing is still cheaper than one year of early replacement. For all assets, four years of postponement is more expensive than one year of early replacement.

Now let us see how the solution planning changes when we implement this non-linear cost structure. We see that the solution planning is now very different from the original solution. For Scenario 1, now 13 assets are replaced early, compared to 14 in the original solution. 55 assets are replaced late, compared to only 9 in the original solution. For Scenario 2, only 13 assets are replaced early, whereas in the original model 58 assets were replaced early. 52 assets are replaced late, compared to 14 in the original model.

Scenario	Experiment	Early replacements	Average years early	Late replacements	Average years late
1	Original	14	1.07	9	1.22
1	5	13	1	55	1.05
2	Original	58	1.12	14	1.29
2	5	13	1.62	52	1.73

Table 21. The results for experiment 5 in terms of early and late replacements.

Interesting to see here is that the number of early and late replacements for Scenario 1 and Scenario 2 is similar. Still, in Scenario 1 there are 25 clusters, whereas in Scenario 2 we only have 12 clusters. In Table 21 we can see however that in Scenario 2 replacements are shifted further, i.e. with more years. When the balancing parameter is still relatively low (in Scenario 1) the model prefers to replace assets only one year late multiple times in the horizon over replacing the assets with more years late in the first replacement. This is different from when we assumed the costs for late replacement to be linear. Scenario 1 therefore results in a very different solution planning. When we increase the balancing parameter however, i.e. when clustering becomes more beneficial, we again see that all shifts are made in the first replacement so that replacement moments are already synchronized as early is possible in the horizon. Scenario 2 is therefore similar to the original situation. Although assets are more often replaced late and therefore replaced in different years, the model also finds a repetitive pattern directly in the beginning of the horizon in which an cluster is planned every five years.

The penalty costs for Scenario 1 are now €46,424. This is 79.5% lower than in the original model where decreasing the number of clusters to 25 resulted in €226,802. Scenario 2 now results in penalty costs of €610,435. This is 42.4% lower than in the original model, when decreasing the number of clusters to 12 costed €1,049,135. We can conclude that the costs functions are very important for the total penalty costs and therefore for the decisions that the model makes. It is therefore recommended to research the financial effects of postponing replacements and refine the costs functions of the model.

5.5. Conclusion

In this chapter we discussed the results of the model for different values of the balancing parameter in Section 5.2 and the values of the other input parameters in Section 5.4. In this section we conclude on our findings.

- *By making relatively small shifts of only one or two years, the number of clusters can already heavily be reduced.*
  Reducing the number of clusters to 25 and 12 requires to shift in total 23 assets and 72 assets respectively, but the number of years with which assets are shifted is limited to only one or two years.

- *Clustering as much as and as early as possible in the horizon is beneficial on the long-term.*
  In order to come to a repetitive pattern or cadence, it is beneficial to make sure to cluster replacements as much as possible and as early in the horizon as possible. We already saw this in Section 5.2 when we found that the vast percentage of the penalty costs result from shifts in the first replacement moment, and only a small percentage comes from shifts in the subsequent replacements. This was confirmed by the experiment in Section 5.3.1, when we limited the values for AE_a and AL_a to 1. This resulted in much higher penalty costs for both scenarios. In the experiment of Section 5.3.2 we assigned a random first replacement F_a between 2019 and 2024 to already fully depreciated assets, instead of setting this first replacement to be in 2019. We have seen that this not only leads to higher penalty costs, but that it also takes more time for the planning to find repetitive cycles with equal intervals.

- *Varying the replacement value for the constructive elements does not influence the solution planning.*
  When varying the replacement values for the constructive elements (since we are less certain about these costs) between 80% and 120% of the original estimated values, we found that the solution planning does not change.

- *Varying the balance between the penalty costs for an early replacement CE_a and the penalty costs for a late replacement CL_a does not influence the solution planning.*
  Only when the costs for one year of late replacement CL_a is set to be larger than $2 \times CE_a$, i.e. two years of late replacement are more expensive than one year of early replacement, the solution starts to change.

- *Implementing a non-linear cost function for late replacements changes the solution planning significantly.*
  This results in more assets being replaced late, and also in more assets being shifted with more than one year (especially in Scenario 2). For Scenario 1, when the balancing parameter is still relatively low, the model now also shifts the subsequent replacements in the horizon instead of only the first, since the costs now increase non-linear. When the balancing parameter is increased all shifts are again made in the first replacement and moments are again synchronized as early in the horizon as possible. Moreover, the cost function that is used in the model is very important for the total penalty costs and therefore for the decisions that the model makes. It is recommended to get a better insight in the financial effects of postponing replacements and to refine the costs functions of the model.

6. Implementation

In this research we developed a model that can help ASM in planning replacements over a long time horizon. Planning over a longer horizon offers opportunities for the clustering of activities, which in turn gives ASM the possibility to provide the main contractor long in advance with time windows in which large-scale renovations and replacements are planned. In this new methodology, it will be the contractors' responsibility to make sure that the assets meet the performance levels as agreed upon, by performing regular maintenance in between these clustered moments. This would allow ASM to grow towards a more controlling role in the partnership.

Since in the new methodology the major maintenance moments with a high impact will be known long in advance, there will be more time to prepare for the execution of these clusters, which should increase predictability and the realization of plans. ASM thinks that these clustered moments create a stage for developing and carrying out improvement projects as opposed to merely replacing individual assets. These will be the moments in which the demands and wishes of the stakeholders can be granted. ASM hopes that this results in stakeholders like Operations or airlines more easily accepting major maintenance projects.

For the implementation of the model, we want to stress that it is highly recommended to use this model in a rolling horizon setting. This means that the model is initially solved with the information that is available at that moment, i.e. average lifecycles. Over time more accurate information about the actual performance and degradation patterns may become available and the expected lifecycle of a specific asset may turn out to fall below or exceed the expected lifecycle as was used in the initial model. In this case, ASM is advised to include this short-term information on the actual performance of the asset in the model, update the input by adjusting the preliminary expected lifecycle and solve the model again. The solution, i.e. the long-term replacement planning, may therefore change when better, short-term information becomes available. This way, the planning becomes more dynamic and accurate. Moreover, it is advised to solve the model again every time replacements are executed, since this changes the construction year (i.e. latest replacement year) and consequently the initial first replacement moment in the horizon. When ASM decided to replace the asset in another year than as was planned by the model, the optimal solution may change. The model is therefore best used in an iterative process as opposed to delivering a one-time, static solution.

A shift towards this new methodology requires however more than only a model. Implementing this model and new methodology is most of all a cultural change. It requires ASM to change the way of thinking and adopt a different approach towards the management of AAS's assets. As we saw in Chapter 2 already, decisions are currently made on asset level. We saw that this leads to a high number of clusters with replacements largely dispersed over time. ASM has to recognize that what is optimal for an individual asset is not necessarily optimal for the area or a process as a whole. This requires a long-term vision, whereas the current decision-making process is mainly ad-hoc. Moreover, adopting the new approach requires a helicopter-view. ASM should start to think more in processes, instead of in individual assets. Clustering replacements implies that assets will not always be optimally utilized. ASM should be willing to accept this and recognize the value that flows from clustering. This value does not only flow from the lower impact on operations but also from the decreased time and effort that is needed to plan the replacements and expected decreased overheads. While in the current situation, all assets are constantly monitored and inspected, in the new methodology ASM is better able to focus on a single area and those assets that are planned for the next cluster. This allows for a more effective and efficient decision-making process. Other benefits result from the fact that in the new methodology, the scope, the scale and the operational impact of the replacements is known longer in advance. This makes the replacements easier to integrate in the daily operational processes and with development projects. Moreover, it offers a stage for developing and carrying out improvement projects as opposed to merely replacing individual assets, and for a better fit with Schiphol's strategic goals.

Quantifying the benefits of clustering is not only helpful for recognizing and communicating the value of a long-term replacement planning, but also when making a reasoned trade-off between the penalty costs that need to be made in order to decrease the number of clusters. The more penalty costs ASM chooses to accept, the more the number of clusters can be decreased. As we already indicated in Chapter 5, how many penalty costs to accept depends on the value that ASM believes to flow from clustering, as well as on the preferences of the different stakeholders. A more in-depth research on the costs that result from the high dispersion of replacements over time in the current situation and the value that is generated by decreasing the number of clusters may help ASM to determine how the penalty costs offset the resulting decrease in the number of clusters. The needs and preferences may differ per stakeholder. Clustering decreases the frequency of maintenance being performed in a certain area, but at the same time the scale of the clusters increases. Some stakeholders may prefer to have a higher number of clusters, but with a smaller impact on operations per cluster. By varying the balancing parameter, ASM is able to steer the model in a certain direction, with more focus on costs or more focus on the number of clusters.

7. Conclusion and recommendations

In this final chapter we will answer the main research question in Section 7.1, whereas Section 7.2 offers several recommendations for the implementation of the new methodology and directions for further development of the model.

7.1. Conclusion

The Asset Management division of Royal Schiphol Group wants to make a shift from planning replacements for individual asset towards an more integrated approach in which replacements are planned over longer time horizons. This long-term view and integrated approach offer potential for decreasing the dispersion of activities over time by clustering replacements. In this research we therefore aimed to answer the following research question:

> *How can the Asset Management division of Royal Schiphol Group plan the replacements of assets over a 60 year horizon, in order to limit the impact on operations?*

We can draw several conclusions that together answer this question:

- *Formulating the problem as an Integer Linear Programming model allows us to quickly find the optimal combination of replacements.*
  This is the combination that minimizes the number of clusters at the lowest cost.

- *What savings in the number of clusters can be achieved depends on how many penalty costs ASM chooses to accept.*
  There is a logarithmic correlation between the penalty costs that are accepted and the savings in the number of clusters that can be achieved. For this E-pier case we found that a decrease in the number of clusters from 51 to 25 can be reached when ASM accepts €226,802 in penalty costs, whereas in order to limit the number of clusters to 12 ASM has to accept penalty costs of €1,049,135.

- *The number of clusters can already be heavily reduced by shifting replacements by only one or two years.*
  These savings in the number of clusters can also be reduced by mainly replacing assets earlier and postponing relatively few replacements. This implies that the impact on operation can be heavily reduced without having to accept increased risks.

- *Aligning replacements as early as possible in the horizon is beneficial over the long term.*
  Synchronizing lifecycles early in the horizon allows us to minimize the number of clusters for relatively low costs, since no additional shifts in subsequent replacements need to be made in order to plan replacements in the same year.

7.2. Recommendations

We will close this chapter with some recommendations related to the use of the model and on which themes to focus when further developing the model.

- *Cluster replacements as early in the horizon as possible.*
  One of the most important recommendations relates to the conclusion in Chapter 5, where we saw that it is beneficial to cluster replacements as much as possible and as early in the horizon as possible. We therefore recommend ASM that when making the transition from the current methodology towards a long-term replacement plan, to invest in rapidly aligning the replacement moments of the assets.

- *Invest in complementing the dataset and enriching the input data.*
  The dataset used for this case includes information of 1,034 in the E-pier, which means the dataset is not yet complete. It is therefore recommended to complement the missing data. Moreover, assumptions for the replacement values have been made. It is recommended to improve the input with more accurate costs parameters, especially for the cost for late replacement. This cost function heavily influences the solution that the model. Furthermore, for this research it was determined to differentiate between asset types and construction years and not to take into account lower levels of detail. This means that assets of the same type were assigned the same values for the input parameters. In the future, ASM may want to add more asset-specific information, such as the replacement value for each individual asset. It should however be questioned if the additional time investment pays off in terms of a higher accuracy in the planning. Moreover, ASM may want to investigate how the costs of late replacement increase over time and refine the cost function in the model accordingly.

- *Investigate the consequences of the apparent gap between depreciation periods and the technical lifetime of assets.*
  In Chapter 2 we concluded that around 30% of the assets in the MC2019 database is already fully depreciated. It should be mentioned here that the condition of most assets is nevertheless good. ASM may however want to reassess the accuracy of the used depreciation periods.

- *Add location data and specific set-up costs.*
  At the moment, data related to the precise location of an asset is limited. Therefore, in this research the possible interaction between assets that share the same location at the E-pier is not taken into account yet. In reality, replacing the lighting and sprinklers at the same gate in the same year may come with additional, financial benefits than is represented by the number of clusters only. This can be modelled by if-then constraints or by taking into account asset-specific set-up costs. Moreover, some replacements may exclude each other, i.e. they should not be planned in the same year. Also dependencies like these can easily be added, by for example either-or constraints. We recommend ASM to invest in increasing the precision and availability of location-related asset data and in addition map how various assets interact. Ideally, in the future the asset data, including information on interactions, can be directly exported from SGIS (Schiphol Geographic Information System) by selecting an area at the terminal complex, such that the model can be solved for that specific area.

- *Extent the model with the planning of major maintenance activities or combine it with the project planning.*
  Lastly, for this research we focused on replacements only. The model can be extended by also taking into account major maintenance moments. Moreover, in the future ASM might want to combine the planning of these maintenance moments and replacements with development projects.

8. References

Budai-Balke. (2009). *Operations Research Models for Scheduling Railway Infrastructure Maintenance.* Erasmus University Rotterdam,

Cho, & Parlar. (1991). A survey of maintenance models for multi-unit systems. *European Journal of Operational Research, 51*, 1-23.

Dekker. (1996). Applications of maintenance optimization models: a review and analysis. *Reliability Engineering and System Safety, 51*, 229-240.

Dekker, Smit, & Losekoot. (1992). Combining maintenance activities in an operational planning phase: a set-partitioning approach. *Journal of Mathematics Applied in Business & Industry, 3*, 315-331.

Dekker, Wildeman, & Van der Duyn Schouten. (1997). A Review of Multi-Component Maintenance Models with Economic Dependence. *Mathematical Methods of Operations Research, 45*, 411-435.

Gits. (1992). Design of maintenance concepts. *International Journal of Production Economics, 24*, 217-226.

Liang. (1985). Optimum Piggyback Preventive Maintenance Policies. *IEEE Transactions on Reliability, 34*(5), 529-538.

Royal Schiphol Group. (2017a). *Accounting Manual*. Retrieved from

Royal Schiphol Group. (2017b). *Annual Report 2017*. Retrieved from

Van Dijkhuizen, & Van Harten. (1997). Optimal clusstering of frequency-constrained maintenance jobs with shared set-ups. *European Journal of Operational Research, 99*, 552-564.

Wildeman, Dekker, & Smit. (1997). A dynamic policy for grouping maintenance activities. *European Journal of Operational Research, 99*, 530-551.

Winston. (2004). *Operations Research: Applications and Algorithms*: Thomson.

Appendix A: Asset database E-pier

Confidential

Appendix B: Model formulation in AIMMS

```
Set Assets {
      Index: a;
      Definition: elementrange(1,116,1,"a-");
   }
   Set Time {
      Index: t, t1;
      Definition: elementrange(1,PlanningHorizon,1,"t-");
   }
   Parameter Lifecycle {
      IndexDomain: a;
   }
   Parameter PlanningHorizon {
      Definition: 60;
   }
   Parameter BalancingParameter;
   Parameter BigM {
      Definition: 120;
   }
   Parameter FirstMoment {
      IndexDomain: a;
   }
   Parameter Allowed_early {
      IndexDomain: a;
   }
   Parameter Allowed_late {
      IndexDomain: a;
   }
   Parameter MinimumInterval {
      IndexDomain: a;
      Definition: Lifecycle(a) - Allowed_early(a);
   }
   Parameter MaximumInterval {
      IndexDomain: a;
      Definition: Lifecycle(a) + Allowed_late(a);
   }
   Parameter Cost_early {
      IndexDomain: a;
   }
   Parameter Cost_late {
      IndexDomain: a;
   }
   Variable x {
      IndexDomain: (a,t);
      Range: binary;
   }
   Variable y {
      IndexDomain: t;
      Range: binary;
   }
   Variable EarlyIndicator {
      IndexDomain: (a,t);
```

```
      Range: binary;
   }
   Variable EarlyIndicatorFirst {
      IndexDomain: (a,t);
      Range: binary;
   }
   Variable LateIndicator {
      IndexDomain: (a,t);
      Range: binary;
   }
   Variable LateIndicatorFirst {
      IndexDomain: (a,t);
      Range: binary;
   }
   Variable PenaltyCosts {
      Range: free;
      Definition: sum[(a,t),EarlyIndicator(a,t)*Cost_early(a)+LateIndicator(a,t)*Cost_late(a)];
   }
   Variable PenaltyCostsFirst {
      Range: free;
      Definition: {
         sum[(a,t)|ord(t)<FirstMoment(a),EarlyIndicatorFirst(a,t)*Cost_early(a)*(FirstMoment(a)-
ord(t))] +

sum[(a,t)|(FirstMoment(a)<ord(t)<=FirstMoment(a)+Allowed_late(a)),LateIndicatorFirst(a,t)*C
ost_late(a)*(ord(t)-FirstMoment(a))]
      }
   }
   Variable BP_Fragmentation {
      Range: free;
      Definition: BalancingParameter * sum[t, y(t)];
   }
   Variable ObjectiveFunction {
      Range: free;
      Definition: PenaltyCosts + BP_Fragmentation + PenaltyCostsFirst;
   }
   Constraint FragmentationConstraint {
      IndexDomain: t;
      Definition: BigM * y(t) >= sum[a,x(a,t)];
   }
   Constraint FirstReplacement {
      IndexDomain: a;
      Definition: sum[t| (FirstMoment(a)-Allowed_early(a) <= ord(t) <=
FirstMoment(a)+Allowed_late(a)),x(a,t)] >= 1;
   }
   Constraint MinimumConstraint {
      IndexDomain: (a,t);
      Definition: sum [t1 | (ord(t1) <= (ord(t) + MinimumInterval(a)-1) and t1 >= t), x(a,t1)] <= 1;
   }
   Constraint MaximumConstraint {
      IndexDomain: (a,t) | ord(t)<(PlanningHorizon-MaximumInterval(a)+2);
      Definition: sum [t1 | (ord(t1) <= (ord(t) + MaximumInterval(a)-1) and t1 >= t), x(a,t1)] >= 1;
   }
   Constraint EarlyCounter {
```

```
      IndexDomain: (a,t);
      Definition: sum [t1 | (ord(t1) <= (ord(t) + Lifecycle(a)-1) and t1 >= t), x(a,t1)] <= (1 + BigM *
EarlyIndicator(a,t));
   }
   Constraint LateCounter {
      IndexDomain: (a,t) | ord(t) < (PlanningHorizon-Lifecycle(a)+2);
      Definition: sum [t1 | (ord(t1) <= (ord(t) + Lifecycle(a)-1) and t1 >= t), x(a,t1)] >= (1 - BigM *
LateIndicator(a,t));
   }
   Constraint FirstReplacementEarly {
      IndexDomain: (a,t) | ord(t) < FirstMoment(a);
      Definition: sum [t1 |ord(t1) <= ord(t) and t1 >= t, x(a,t)] <= (BigM * EarlyIndicatorFirst(a,t));
   }
   Constraint FirstReplacementLate {
      IndexDomain: (a,t) | FirstMoment(a) < ord(t) <= FirstMoment(a)+Allowed_late(a);
      Definition: sum [t1 |ord(t1) <= ord(t) and t1 >= t, x(a,t)] <= (BigM * LateIndicatorFirst(a,t));
   }
   MathematicalProgram LongTermMaintenancePlanning {
      Objective: ObjectiveFunction;
      Direction: minimize;
      Constraints: AllConstraints;
      Variables: AllVariables;
      Type: Automatic;
   }
}
```

Appendix C: Estimating replacements values

Confidential

Appendix D: Sensitivity analysis, experiment 3

Scenario	**Experiment**	**Description**	**Penalty costs**
1	Original	Original replacement value	€ 226,802
1	3a	Replacement value = 0,8 × original replacement value	€ 226,682
1	3b	Replacement value = 0,9 × original replacement value	€ 226,742
1	3c	Replacement value = 1,1 × original replacement value	€ 226,862
1	3d	Replacement value = 1,2 × original replacement value	€ 226,922
2	Original	Original replacement value	€ 1,049,135
2	3a	Replacement value = 0,8 × original replacement value	€ 1,005,887
2	3b	Replacement value = 0,9 × original replacement value	€ 1,027,511
2	3c	Replacement value = 1,1 × original replacement value	€ 1,070,760
2	3d	Replacement value = 1,2 × original replacement value	€ 1,092,384

Table 1.The penalty costs for the modified replacement values for constructive assets.

Appendix E: Modified model with non-linear costs for late replacement

Mathematical formulation

Sets

A	set of assets a
T	years $t \in \{1 \dots T\}$
U	time units u in years

Parameters

PH	planning horizon
B	balancing parameter that balances the importance of clustering as opposed to the penalty costs
LC_a	the lifecycle of asset a
AE_a	the allowable number of years that the replacement of asset a may be executed early
AL_a	the allowable number of years that the replacement of asset a may be executed late
Min_a	$LC_a - AE_a$, the minimum interval between two replacements of asset a
Max_a	$LC_a + AL_a$, the maximum interval between two replacements of asset a
F_a	the original first execution moment of the replacement of asset a in the time horizon, without shifting
CE_a	the penalty cost for replacing asset a one year earlier than at its end-of-life (linear in the number of years the replacement is planned early)
CL_a	the penalty cost for replacing asset a one year later than at its end-of-life (linear in the number of years the replacement is planned late)

Decision variables

$x_{a,t}$	binary variable that denotes if a replacement of asset a is planned in year t (1) or not (0)
y_t	binary variable that denotes if at least one replacement is planned in year t ($y_t = 1$ if $\sum_a x_{a,t} \geq 1$; 0 if not)
$EF_{a,t}$	binary variable that denotes if the first replacement of asset a is executed early in year t
$LF_{a,t}$	binary variable that denotes if the first replacement of asset a is executed late in year t

$E_{a,t}$ binary variable that denotes if a replacement of asset a at time t is followed by an early replacement (1) or not (0)

$L_{a,t,u}$ binary variable that denotes if a replacement of asset a at time t is followed by a late replacement of at least u units (1) or not (0)

$$Min\left(B\sum_t y_t + \sum_a \sum_t^{t=F_a-1} (EF_{a,t} * CE_a * (F_a - t)) + \sum_a \sum_{t=F_a+1}^{t=F_a+AL_a} (LF_{a,t} * CL_a * (t - F_a)^2) + \sum_a \sum_t (E_{a,t} * CE_a) + \sum_a \sum_t \sum_u (L_{a,t,u} * CL_a * u) \right)$$

s.t.

$$\sum_a x_{a,t} \leq My_t \qquad \forall t \qquad (1)$$

$$\sum_{t=F_a-AE_a}^{F_a+AL_a} x_{a,t} = 1 \qquad \forall a \qquad (2)$$

$$x_{a,t} \leq M * EF_{a,t} \qquad \forall a, t < F_a \qquad (3)$$

$$x_{a,t} \leq M * LF_{a,t} \qquad \forall a, F_a < t \leq F_a + AL_a \qquad (4)$$

$$x_{a,t} + \cdots + x_{a,t+Min_a-1} \leq 1 \qquad \forall a, \forall t \qquad (5)$$

$$x_{a,t} + \cdots + x_{a,t+Max_a-1} \geq 1 \qquad \forall a, \forall t \qquad (6)$$

$$x_{a,t} + \cdots + x_{a,t+LC_a-1} \leq 1 + M * E_{a,t} \qquad \forall a, \forall t \qquad (7)$$

$$x_{a,t} + \cdots + x_{a,t+LC_a-1+(u-1)} \geq 1 - M * L_{a,t-1,u} \qquad \forall a, t < PH - Max_a + 2 \qquad (8)$$

Model formulation in AIMMS

```
Set Assets {
    Index: a;
    Definition: elementrange(1,116,1,"a-");
  }
  Set Time {
    Index: t, t1;
    Definition: elementrange(1,PlanningHorizon,1,"t-");
  }
  Set Years_u {
    Index: u, u1;
    Definition: elementrange(1,10,1,"u-");
  }
  Parameter Lifecycle {
    IndexDomain: a;
  }
  Parameter PlanningHorizon {
    Definition: 60;
  }
  Parameter BalancingParameter;
  Parameter BigM {
    Definition: 120;
  }
```

```
  Parameter FirstMoment {
    IndexDomain: a;
  }
  Parameter Allowed_early {
    IndexDomain: a;
  }
  Parameter Allowed_late {
    IndexDomain: a;
  }
  Parameter MinimumInterval {
    IndexDomain: a;
    Definition: Lifecycle(a) - Allowed_early(a);
  }
  Parameter MaximumInterval {
    IndexDomain: a;
    Definition: Lifecycle(a) + Allowed_late(a);
  }
  Parameter Cost_early {
    IndexDomain: a;
  }
  Parameter Cost_late {
    IndexDomain: a;
  }
  Variable x {
    IndexDomain: (a,t);
    Range: binary;
  }
  Variable y {
    IndexDomain: t;
    Range: binary;
  }
  Variable EarlyIndicator {
    IndexDomain: (a,t);
    Range: binary;
  }
  Variable EarlyIndicatorFirst {
    IndexDomain: (a,t);
    Range: binary;
  }
  Variable LateIndicator {
    IndexDomain: (a,t,u);
    Range: binary;
  }
  Variable LateIndicatorFirst {
    IndexDomain: (a,t);
    Range: binary;
  }
  Variable PenaltyCostsEarly {
    Range: free;
    Definition: sum[(a,t),EarlyIndicator(a,t)*Cost_early(a)];
    Comment: "Oude functie:
sum[(a,t),EarlyIndicator(a,t)*Cost_early(a)+LateIndicator(a,t)*Cost_late(a)]";
  }
  Variable TotalPenaltyCostsLate {
```

```
        Range: free;
        Definition: sum[(a,t,u),PenaltyCostsLate(a,t,u)];
        Comment: "PenaltyLate(a,t,u) - sum[u1 | ord(u1) < ord(u) , PenaltyLate(a,t,u1)]";
    }
    Variable PenaltyCostsLate {
        IndexDomain: (a,t,u);
        Definition: Cost_late(a) * LateIndicator(a,t,u) * ord(u);
    }
    Variable PenaltyCostsFirst {
        Range: free;
        Definition: {
            sum[(a,t)|ord(t)<FirstMoment(a),EarlyIndicatorFirst(a,t)*Cost_early(a)*(FirstMoment(a)-
ord(t))] +

sum[(a,t)|(FirstMoment(a)<ord(t)<=FirstMoment(a)+Allowed_late(a)),LateIndicatorFirst(a,t)*C
ost_late(a)*((ord(t)-FirstMoment(a))^2)]
        }
    }
    Variable BP_Fragmentation {
        Range: free;
        Definition: BalancingParameter * sum[t, y(t)];
    }
    Variable ObjectiveFunction {
        Range: free;
        Definition: BP_Fragmentation + PenaltyCostsFirst + PenaltyCostsEarly +
TotalPenaltyCostsLate;
    }
    Constraint FragmentationConstraint {
        IndexDomain: t;
        Definition: BigM * y(t) >= sum[a,x(a,t)];
    }
    Constraint FirstReplacement {
        IndexDomain: a;
        Definition: sum[t| (FirstMoment(a)-Allowed_early(a) <= ord(t) <=
FirstMoment(a)+Allowed_late(a)),x(a,t)] >= 1;
    }
    Constraint MinimumConstraint {
        IndexDomain: (a,t);
        Definition: sum [t1 | (ord(t1) <= (ord(t) + MinimumInterval(a)-1) and t1 >= t), x(a,t1)] <= 1;
    }
    Constraint MaximumConstraint {
        IndexDomain: (a,t) | ord(t)<(PlanningHorizon-MaximumInterval(a)+2);
        Definition: sum [t1 | (ord(t1) <= (ord(t) + MaximumInterval(a)-1) and t1 >= t), x(a,t1)] >= 1;
    }
    Constraint EarlyCounter {
        IndexDomain: (a,t);
        Definition: sum [t1 | (ord(t1) <= (ord(t) + Lifecycle(a)-1) and t1 >= t), x(a,t1)] <= (1 + BigM *
EarlyIndicator(a,t));
    }
    Constraint LateCounter {
        IndexDomain: (a,t,u) | ord(t) < (PlanningHorizon-Lifecycle(a)+2);
        Definition: sum [t1 | (ord(t1) <= (ord(t) + Lifecycle(a)-1 + ord(u-1)) and t1 >= t), x(a,t1)] >=
(1 - BigM * LateIndicator(a,t-1,u));
```

```
    Comment: "sum [t1 | (ord(t1) >= ord(t) and t1 <= t + (ord(u)-1)), LateIndicator(a,t1)] <=
(ord(u)-1) + 100 * NIEUW_NumberLate(a,t,u)";
  }
  Constraint FirstReplacementEarly {
    IndexDomain: (a,t) | ord(t) < FirstMoment(a);
    Definition: sum [t1 |ord(t1) <= ord(t) and t1 >= t, x(a,t)] <= (BigM * EarlyIndicatorFirst(a,t));
  }
  Constraint FirstReplacementLate {
    IndexDomain: (a,t) | FirstMoment(a) < ord(t) <= FirstMoment(a)+Allowed_late(a);
    Definition: sum [t1 |ord(t1) <= ord(t) and t1 >= t, x(a,t)] <= (BigM * LateIndicatorFirst(a,t));
  }
  MathematicalProgram LongTermMaintenancePlanning {
    Objective: ObjectiveFunction;
    Direction: minimize;
    Constraints: AllConstraints;
    Variables: AllVariables;
    Type: Automatic;
  }
```