Production-scheduling in Fresh Vegetable Processing

A case study at W. Heemskerk BV.

MSc Thesis

Author:
Cornelis Petrus van Veen (362926)

Supervisors:
Dr. W. van den Heuvel (EUR)
S. Grimbergen (Heemskerk)

ERASMUS UNIVERSITEIT ROTTERDAM

Heemskerk
fresh & easy

Rijnsburg, November, 2013

Preface

The completion of this thesis marks both the end of my studying time and my participation in the master program Operations Research and Quantitative Logistics at the Erasmus University in Rotterdam. Although the route has not always been straightforward, I firmly believe that the completion of this thesis is the crown on a very interesting and rewarding period in my life. Although I now only hold a small foothold in the vast scientific world, I was given the knowledge and resources to further extend my reach.

This thesis is the result of a 6-month-long internship at W. Heemskerk BV I would like to thank Dr. Wilco van den Heuvel for his continuous support and availability throughout the course of this thesis. His insights were able to get and keep me on track during the entire period.

Furthermore I would like to extend my gratitude to W. Heemskerk BV and especially to Sonja Grimbergen and Remon Post for providing me with the opportunity to investigate this interesting problem and providing me all necessary resources at their company. I would also like to thank all my colleagues at W. Heemskerk BV for making available their time and resources to provide me with the required data and their useful insights.

Last, but not least, I would like to thank my closest relatives who have given me unconditional support throughout the course of my study.

Finally, I hope that this thesis provides useful insights for others and that it sparks the same enthusiasm for the subject as I have.

Notation

Sets

W	The set of batches (work orders)
$W^m \subset W$	The (sub)set of make batches (batches that produce intermediate products)
$W^p \subset W$	The (sub)set of pack batches (batches that produce final products)
$W_w^m \subseteq W^m$	The (sub)set of make batches that supply pack batch $w \in W^p$
$W_w^r \subseteq W^m$	The (sub)set of make batches required to be finished before make batch $w \in W^m$ can be started
J	The set of work centers (production units)
$J_w \subseteq J$	The (sub)set of work centers that can produce batch $w \in W$
$J_{ww'} \subseteq J$	The (sub)set of work centers that can produce both $w \in W\ and\ w' \in W$

Parameters

α_w	Equal to the processing time, in hours, required to produce one unit of batch $w \in W$
β_w	Number of units required to be produced of batch $w \in W$
D	Preferred time, in hours, at which all work centers should have completed their batches
D_w	Preferred time, in hours at which work order $w \in W$ should be finished
δ_w	Equal to the time required before batch $w \ \in W$ can be used in another batch
G_w	This parameter contains the starting times of work orders $w \in W$
$K_{ww'}$	Equal to 1 if for a pair of batches w and w' the variable $Y_{ww'}$ should be fixed to 0, 0 otherwise.
M	A sufficiently large number
$\omega_{ww'}$	A penalty for a switchover from batch $w \in W$ to batch $w' \in W$

Variables

F_{wj} Binary variable. Equal to 1 if batch $w \in W$ starts the processing sequence on workcenter $j \in J_w$, 0 otherwise.

X_{wj} Binary Variable. Equal to 1 if batch $w \in W$ is assigned to work center $j \in J$ but not the first to be processed, 0 otherwise.

$Y_{ww'}$ Binary Variable. Equal to 1 if batch $w \in W$ is processed right before batch $w' \in W$, 0 otherwise

S_w Nonnegative variable. Equals the start time, in hours, of batch $w \in W$

D_w^+ Nonnegative variable. Equal to the amount of time that batch $w \in W^p$ finishes after the preferred completion time D

Abstract

W. Heemskerk BV is one of the largest fresh processed vegetable production companies in Europe. Weekly, more than two million products find their way to well-known supermarket chains and fast-food restaurants in the Netherlands and abroad.

The production process consists of two well-defined stages. In the first stage, vegetables are cut and washed into a large range of intermediate products. In the second stage, these products are packaged into various types of packaging material. Most of this work is performed on highly specialized production lines.

The processed vegetables have a very short lifetime. This means that the production follows a daily cycle. Each day the entire range of products is produced at least once. In the current situation the production schedule is fixed. Each product is initially assigned to the same machine in the same sequence. To arrive at a feasible production schedule that satisfies all constraints, experienced team leaders manually adjust the schedule during the production day.
This practice of manually adjusting the production schedule is not ideal. Due to the complexity of the process, changes in the schedule can only be made very close to the actual production time. At that point many potentially better options have already been excluded. Furthermore, the production process is constantly becoming more complex. New production lines and product recipes are constantly being added. Because of this, the process is a good candidate for combinatorial optimization.

In this thesis a mathematical formulation is given for the scheduling problem. The production process includes various process-specific constraints. For example, the production sequence of the different products is restricted because of cross-contamination risks. All these process-specific constraints are included in the model. This results in a continuous-time immediate precedence batch-scheduling model that minimizes the amount of bad changeovers, the idle time in the schedule and the tardiness of the different batches.

A first attempt was made to solve the MIP-model by making use of the commercial solver CPLEX. Due to the sheer size of problem instance, this solver is unable to find a feasible solution for the model in a reasonable amount of time. Therefore, a sequential approach is proposed to solve the problem.

The sequential approach consists of a number of steps. In the first step, only the packaging orders are scheduled. The solution of this first step is used as input for the second step in which all pre-processing orders are scheduled. A last step is then required to re-solve the schedule of the packaging orders to arrive at an overall feasible production schedule. For each step of this approach, the same commercial solver is utilized. However, due to the sequential approach, the number of variables and constraints in each of the problem instances is greatly reduced. This results in a greatly reduced solution time.

The sequential approach is compared to the MIP-based approach using a relatively small problem instance. The results of this comparison show that the sequential approach gives solutions of equal or better quality than the MIP-based approach while requiring considerably less solution time.

The capabilities of the sequential approach are further amplified by solving an actual real-world problem instance and by comparing it to the current situation at the case-company.

Key words: Production-scheduling, Optimization, Modeling, Fresh Vegetable Production.

Contents

1. Introduction

The last part of the master program Operations Research and Quantitative Logistics at the Erasmus University is the completion of the master's thesis. This report is the final part of the thesis. In this section, an introduction to production-scheduling problems and the case-company will be given. Furthermore, the contents of the rest of the paper will be introduced.

1.1 Production-scheduling

Production processes have traditionally seen an increase in complexity and a strive for more efficiency. To cope with this, production planning and scheduling has seen a large amount of interest during the past few decades. The problem has been studied extensively by an array of different disciplines. This has resulted in a very large library of solution approaches. Creating a feasible and favorable production schedule has turned out to be very difficult when the number of variables and constraints increases. Many of the problems considered are strongly NP-hard. Herrmann (2006) states that algorithms that can find optimal solutions in polynomial time are unlikely to exist.

The introduction of computers to business' in the 1950s has resulted in a large development of software solutions in the following decades. Today, the majority of production companies use material requirement planning (MRP) in Enterprise Resource Planning (ERP) software. These modules are used to plan material requirements and release purchasing and work orders to the shop floor. Several methods exist to schedule these orders in an adequate way. Some businesses may release the orders to the floor, where experienced workers make a decision about the schedule. Others may use more advanced methods like scheduling algorithms. It is generally believed that much can be improved with regard to manual scheduling. Because of this, academics have been trying to formulate scheduling techniques for decades. Software developers have always been eager to implement these techniques. Solutions incorporating these techniques are often named Advanced Planning and Scheduling (APS) modules (Gayialis & Tatsiopoulos, 2009).

APS modules are usually implemented as extensions to already existing ERP systems. After the orders have been released from the MRP system, they are scheduled in the APS module. Unfortunately, many companies are unsatisfied with their APS implementations (Stadtler & Kilger, 2005). A case study performed by Kjellsdotter (2010) studies the implementation of an APS system in an equipment production company. She finds that general objectives like the minimization of total tardiness or the make span often do not directly apply to (all of) the objectives that a company may have regarding their production scheduling. Another important aspect of the implementations is the acceptance of the generated schedule by the workers on the shop floor. The study shows that when the idea behind the new schedule is not clearly explained, workers tend to doubt the correctness of it and implement their own version of the schedule. This supports the idea that production-scheduling problems are very company specific problems that cannot be caught by a single general model and that a more customized approach should be pursued.

1.2 The case-company

W. Heemskerk BV was founded in 1960. Today, it is one of the top 10 vegetable processing companies in Europe. Heemskerk processes vegetables in Rijnsburg (The Netherlands) and employs approximately 500 employees. Every day, around 60 types of vegetables are processed into as many as 300 different end products. In 2012, more than 80 million fresh products were delivered to well-known supermarkets, fast-food restaurants and catering firms in the Netherlands and abroad.

Heemskerk operates in a complex and demanding market. Customers expect near 100% fill rates and a high flexibility. Furthermore, the availability of unprocessed vegetables varies highly over the seasons. This, along with the fact that the lifetime of products ranges from just 6 to 9 days, makes the production process complex to manage.

In order to cope with high demands of customers and perishability issues, all the different end products are produced every single day. This means the production process is highly flexible and that sufficient capacity is usually available. However, because the lifetime of products is so short, it is not possible to produce products in advance. This means the workload cannot be spread over a month or even over a week. Everything has to be produced one or at most two days before delivery.

The case-company currently employs an ERP system based on MRP. The planning department makes demand prognoses 2 weeks in advance. From these, the total number of required end products, intermediate products and unprocessed vegetables is calculated per day. These figures are then used by the buying department to make sure that enough unprocessed vegetables are available. One day before the actual production day the final demand of end products is determined from sales orders. With this demand, work orders for end products and intermediate products are released to the shop floor for production on the next day.

Production on the shop floor occurs from 6 a.m. until approximately 3 a.m. on the next day. In the current situation, every product is assigned a fixed routing. This means that a certain product will always be planned on the same machine in the same order. After the work orders have been released to the shop floor, an experienced team leader or operator will decide the final machine assignment and production sequence. This usually means that a considerable amount of the work will be produced on a different machine and in a different order than originally planned. The amount of changes required to reach the production goals increases as the workload increases.

The fresh vegetable market is quickly changing and new products are constantly being introduced. This, along with a constant increase in volume, makes production capacity more scarce and the scheduling more complex. As humans are not well equipped to control or optimize a large and complex production schedule by hand, Heemskerk investigated the possible implementation of an APS system. They quickly found that general APS systems are not able to cope with the demands of the company. Managers at Heemskerk do however realize that a lot can be improved by investing in a more sophisticated approach to machine allocation and production scheduling.

Throughout the rest of this thesis W. Heemskerk BV will be mentioned as either ‘Heemskerk’ or the ‘case-company’.

2. Problem Description

In this section the production-scheduling problem at Heemskerk will be explained. This problem may also be referred to as 'the problem' or 'the case' throughout the rest of this thesis. A set of research questions will be given that will provide a guideline for the rest of the thesis.

2.1 Description and research questions

In the current situation at Heemskerk, work orders are released to the shop floor with a fixed machine allocation and production sequence. It is then up to experienced team leaders and operators to adjust this as they see fit. The size and the complexity of the production process make it impossible to do this optimally by hand. In practice, managers see that this results in team leaders optimizing the process locally in their own department or even their own shift. Although the preconditions are met, this never ensures an optimal solution for the entire production chain.

In figure 2.1 the entire production chain of Heemskerk is shown. Unprocessed vegetables and other goods are delivered and stored in the unprocessed goods inventory. From this inventory the goods are released to the different pre-processing production lines when requested by the operators. On these production lines products are washed, cut and mixed to form intermediate products. These intermediate products are then stored into small metal storage containers called 'trolleys'. Next, these trolleys are pulled into the buffer inventory by a conveyor system. Intermediate products can be stored in this buffer inventory for a limited amount of time due to capacity and lifetime restrictions. From the buffer, intermediate products are released to the different packaging lines when requested by the operators of the machines. The packaging lines can be divided into 3 distinct types. Products cannot be interchanged between these lines because of their package type or their specific production process. Distinction is made between lines that package into bags, into bowls and lines where extra hygienic rules apply (so-called HighCare lines). After the packaging is complete, the finished products are stored at the logistics department until their orders are due and the delivery to the customer can be completed.

The current approach to scheduling and machine allocation results in solutions that are at most locally optimal for the pre-processing and packaging stages of the production process. Any infeasibilities (product unavailability, under capacity, etc.) that result from this are solved by either postponing production or by advancing production of intermediate products to the day before they are required. The managers of the company realize that this approach is suboptimal and have requested a solution approach that results in a better connection between the pre-processing and packaging production stages. Of course this approach should not violate any of the constraints set by the surrounding departments. To support this request, a number of research questions will be proposed here. Throughout the following sections an attempt is made to answer these questions.

The above results in the following research question for the thesis:

> ***How can the connection between the pre-processing and packaging production stages at Heemskerk be optimized by the allocation of production resources to the available work?***

To be able to answer this question, at least the following sub-questions should be answered:

- ***How can the problem be modeled using a mathematical formulation?***
  - *Which data will be required in order to correctly model the production process*
  - *What are the constraints that define the production process?*
  - *What are the main objectives of the production process?*
- ***How can we solve this mathematical formulation within a reasonable amount of time?***
- ***How usable is the solution and how does it compare to the current situation?***

These research questions should give the perspective from which the topic of the thesis is approached.

2.2 Problem complexity

Isolating the two High Care packaging lines and the orders that are need to be scheduled on these lines can show the complexity of the problem. This isolated problem of two production lines can be described as a parallel machine-scheduling problem without preemption. Preemption means that the processing of an order can be stopped to rush through another order first.

In a parallel machine-scheduling problem, a set of w orders has to be scheduled on a set of j identical machines. A popular performance measure for the generated schedule is the weighted total tardiness of the orders $\sum p_w D_w^+$. Here p_w is the weight that is given for each order m and D_w^+ is the tardiness of order w. D_w^+ is defined by $D_w^+ = \max\{0, S_w + \alpha_w \beta_w - D_w\}$ where $S_w + \alpha_w \beta_w$ is equal to the finishing time of order w and D_w is equal to the due time of order w.

It was already shown in 1977 that such a parallel machine scheduling problem without preemption and where the weighted total tardiness is minimized belongs to the class of NP-Complete problems even for $j = 2$ (Lenstra & Rinnooy Kan, 1977). The isolated exemplary problem of the two high care lines would thus already be NP-complete.

The minimization of the total tardiness is one of the performance measures that will be included in the model. The two High Care lines only contribute to a very small part of the total problem. The problem presented here has several aspects that add to the complexity of the problem; there is a large amount of unrelated parallel machines, it includes sequence dependent changeovers and a lot of other process-specific aspects.

The presented problem is therefore at least NP-hard. Each day there are more than 700 orders waiting to be produced. This makes the problem a very large NP-hard problem.

Flow Diagram of the production process at Heemskerk

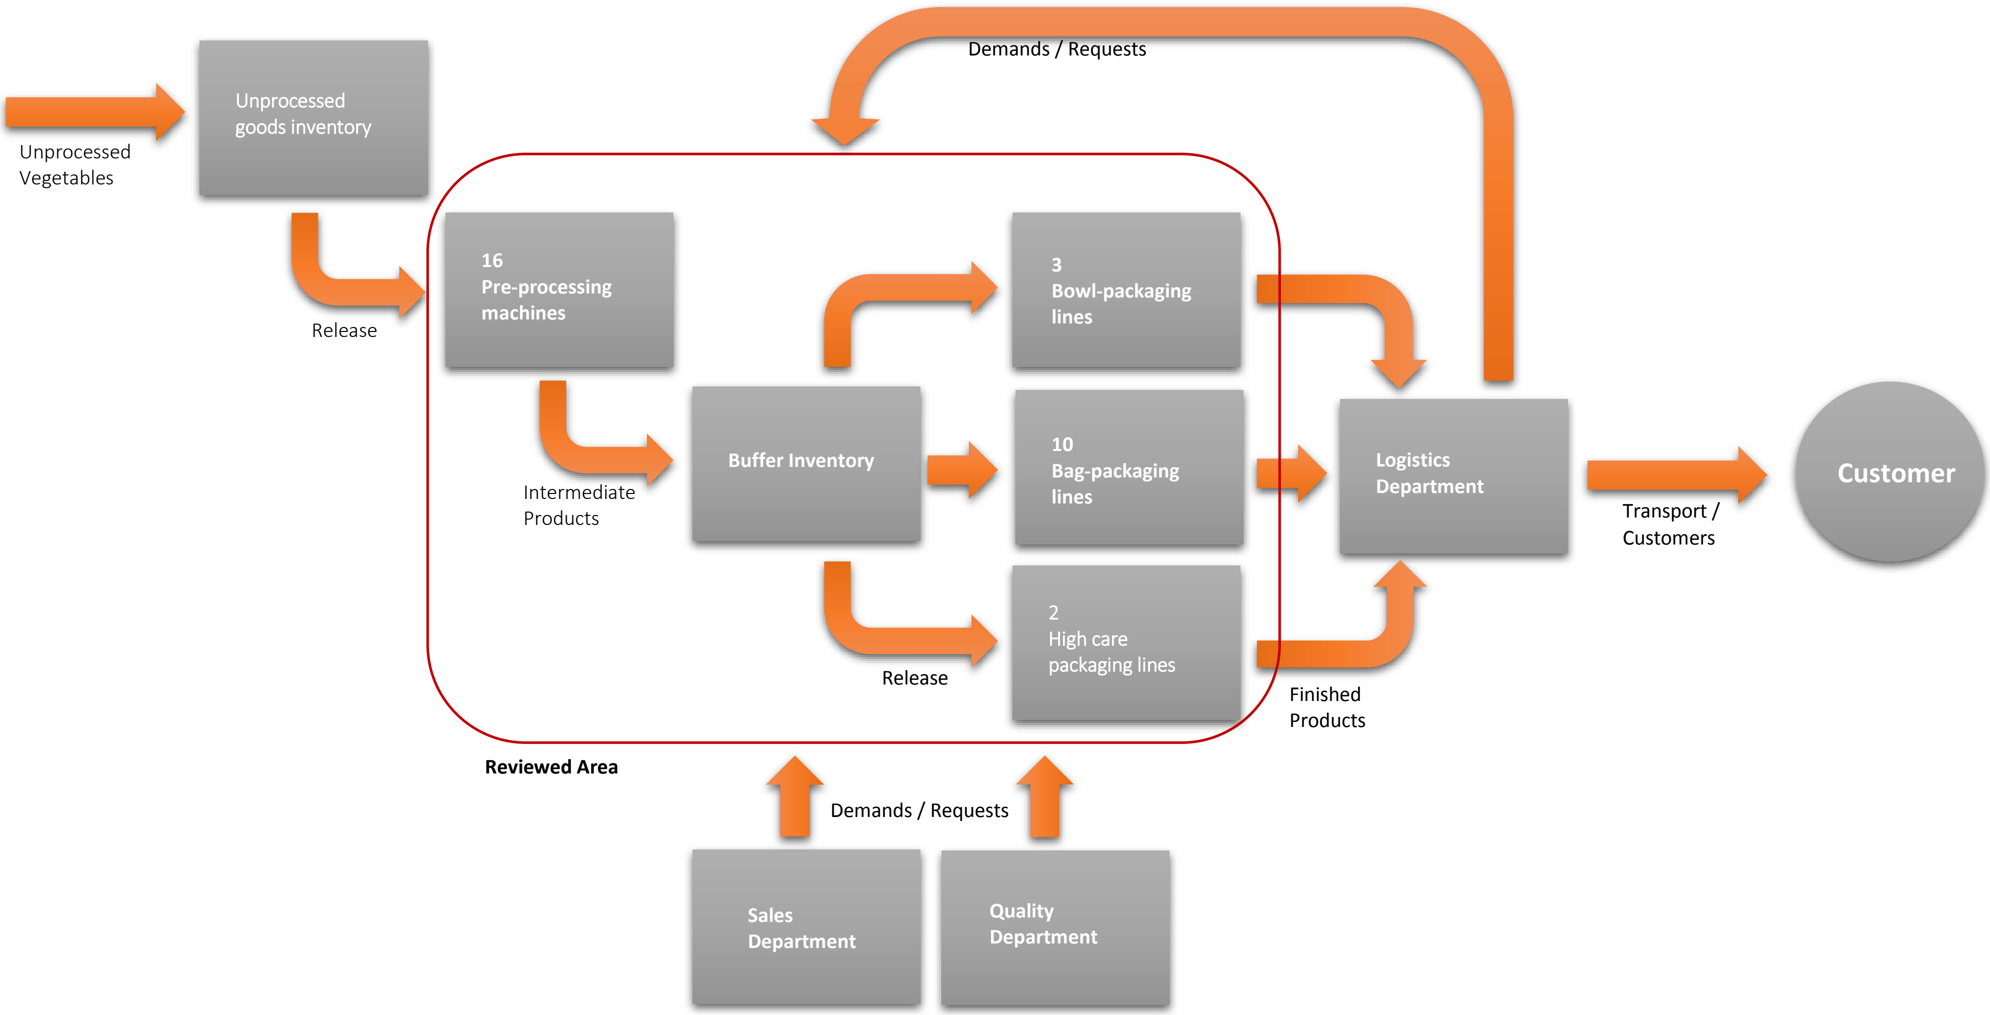

Figuur 2.1: Flow diagram of the production process at Heemskerk

3. Literature Review

In this section a review is presented of literature that was deemed relevant for this research. Different aspects of the problem can be supported by interesting views from other authors. Their findings will support the different choices that were made during the course of this research.
The use of some of the information presented in this section might not directly be apparent. The actual use of this information will be further explained in the relevant sections. Where possible, references will be made to the studies presented in this section.

3.1 APS systems

Because the implementation of an APS system was investigated by the case-company, some research was done into the implementation of APS systems by other companies. Although the interest for APS systems has increased a lot, only a few studies into the user experiences and difficulties exist (Lin, Hwang, & Min-Yang Wang, 2007). Lin, Hwang, & Min-Yang Wang (2007) find that the implementation of an APS system at a semiconductor manufacturer did not result in a reduction of manual labor by the planners. In fact, there was no easy method to verify correctness of the model and the IT specialists or planners could not fully grasp solutions produced by the system. A similar conclusion is drawn by Kjellsdotter (2010) for the implementation of an APS system at an equipment manufacturing company. In addition she finds that the scheduling objectives of the system are hard to match with the actual objectives of the company. Overall, many companies are unsatisfied with their APS implementations (Stadtler & Kilger, 2005). This raises the question of whether the general scheduling models included in APS systems can ever provide satisfactory results.

3.2 Discrete and continuous time representations

The most widely reported methods for process scheduling problems are based on mathematical programming approaches. In recent years, especially methods based on Mixed Integer Linear Programming (MILP) have become popular because of their flexibility and extensive modeling capabilities (Christodoulos, 2005). One of the first major decisions in the formulation of such a model is to choose a representation of time. Time representations can be either discrete or continuous. Stefansson & Sigmarsdottir (2011) give an extensive comparison between discrete and continuous time representations for a large scheduling problem in the pharmaceutical industry. They state that discrete formulations have the advantage that they are in general very flexible and are capable of incorporating many scheduling features. However, two major drawbacks are given. The time domain representation is only an approximation of the real time needed and the number of variables and constraints often become very large when actual industrial problems are being modeled. In addition, it can be difficult to model operations where the processing time is dependent on the batch size. Continuous time representations on the other hand, generally result in a lower number of constraints and variables. The constraints required for this formulation can be more complicated. This can result in a more computationally complex model even though the model has fewer constraints. They conclude that their continuous model results in more useful solutions and generally outperforms the discrete model for the larger problem instances.

3.3 Production-scheduling models

There is a wide range of different production-scheduling models available. Therefore, choosing a model type that fits the problem presented here is no easy task. Each model type has its own strengths, weaknesses and limitations. As was mentioned in the introduction, scheduling problems are some of the hardest combinatorial problems to exist. They belong to the set of NP-complete problems (Garey & Johnson, 1979), which results in solution times that scale exponentially with the size of the problem. This means that choosing the right model for the problem is of vital importance for the project.

Over the years a number of studies comparing different scheduling models were written. Reklaitis (1992) presents an overview of different methods for the scheduling and planning of batch processes. Focus was given to the main components of the chemical process and the available solution methods. Grossmann, Quesada, Raman, & Voudouris (1996) created an overview of mixed-integer optimization techniques. They also include models for the design of chemical batch processes and focus on recent basic solution methods for the problems. Later Pekny & Reklaitis (1998) discussed the different solution techniques available to them. Their study includes an overview of randomized search, constraint guided search, simulation-based strategies and basic mathematical programming approaches.

More recently, a good overview of different MILP-based approaches for the chemical processing industry is given by Floudas & Xiaoxia (2005). This study is focuses on general network represented processes. The choice of time representations is presented as one of the most important decisions. Different models are presented for both discrete- and continuous-time scale representations. Continuous-time models are further split up into Global event-based, Unit-specific event based models. They find that the latter generally results in models of smaller size in terms of the number of variables, compared to discrete-time models and global event-based models. Furthermore, different methods to decrease the time required to solve the models are presented. These methods are supported with examples of small-scale problems.

Another excellent overview of optimization methods for the scheduling of batch processes is presented by Mendez & Cerda (2006). Their review is very complete. In addition to the models presented by Floudas & Xiaoxia (2005) they present continuous-time precedence based models. This includes unit-specific precedence, immediate precedence and general precedence based models. This type of model is very suited for sequential batch oriented processes where changeovers are of concern. General precedence models result in less variables than unit-specific and immediate precedence based models but lack the possibility to model the exact results of sequence-dependent changeover times.

3.4 Solving real-world scheduling problems

Due to the complex nature of scheduling problems, solving actual real-world industrial problems is often challenging. It is now widely recognized that practice and theory should come closer together (Mendez & Cerda, 2006) and (Ruiz & Stutzle, 2008). New academic studies are usually tested on relatively small problems while real-world industrial problems can contain hundreds of batches. A number of studies investigate different methods to solve large-scale problem.

Kopanos, Mendez, & Puigjaner (2010) state that in addition to hundreds of batches, real-world industrial applications have multiple processing units available for each task and their processes are made up from a long sequence of processing stages. Additionally, to create a feasible schedule a wide range of operational constraints has to be formulated. Therefore, optimization solvers have to find optimal or near-optimal solutions in a huge search space within a small feasible region. In their study they introduce a MIP-based solution strategy for a real-world scheduling problem in the pharmaceutical industry. To solve the problem they introduce a 2-step optimization approach where they first insert all orders 1-by-1 and solve the MIP model iteratively. In the second step, after all orders have been inserted, this solution is iteratively improved by taking out and reinserting orders. To quickly solve the MIP model in each iteration, variables of the already inserted orders are fixed to some extent. With this approach they are able to obtain good solutions for a large problem (60 products) within an hour.

Kopanos, Puigjaner, & Geogiadis (2012b) first optimize their model formulation of an ice-cream production process by introducing an additional set of tightening constraints. The model is then further refined by introducing another set of constraints that breaks the symmetry of allocation decisions. To solve the model another MIP-based solution strategy is introduced. This strategy uses specific aspects of the production process to find solutions efficiently. Again, the orders are inserted into the problem iteratively, however this time they are inserted according to their rank. The nature of the production process allows an easy division of the orders into ranks. Products with a higher rank are inserted into the problem before orders with a lower rank. This results in very good solutions without extensive computational requirements.

Wauters & Verbeeck (2012) developed a specialized scheduler for real-world production-scheduling problems. Their approach focuses mainly on finding acceptable, better than manual, solutions quickly. The presented approach uses multiple steps to construct a production schedule. The scheduler uses a routing problem to first assign all orders a route through the production plant. The route that has most capacity left is chosen for each order here. After that a translation step translates the problem into multiple resource-constrained project scheduling problems that are then solved by use of fast heuristics. Finally, an improvement step improves the solution in a guided way. Many optimization methods can serve as a guiding mechanism here. Examples include steepest descent, simulated annealing and tabu-search.

3.5 Planning and scheduling in the processing industry

Although a lot of research has been performed in the field of production scheduling, relatively little attention was given to the food processing industry. Here, an overview is given of literature that is, at least partly, similar to the problem studied for this thesis.
Mendez & Cerdá (2002) present a MILP-based approach for a scheduling problem in a make-and-pack production process. Although their formulation lacks an extensive objective function and the studied problems are of a relatively small size, the structure of the process is similar to the structure of the case-company.

Entrup & Günther (2005) present three MIP models for scheduling and planning problems in yoghurt packaging lines. The process is somewhat similar to fresh vegetable processing. Shelf life issues and intermediate product storage capacities were included, however the models lack details on changeovers and their costs. Marinelli & Nenni (2007) present a model for a similar yoghurt production line and give an effective two-stage heuristic to solve the problem.
Akkerman & Donk (2009) realized that very few studies address the scheduling problem in the food-processing industry and present a general analysis methodology for scheduling problems in food processing.

More recently, a number of contributions to the field were made to reduce the gap between practice and theory in the area of short-term scheduling. Kopanos & Puigjaner (2010) present a mixed discrete/continuous-time MIP formulation for the simultaneous lot sizing and scheduling in yoghurt packing lines. This model is more specialized and now also includes sequence-dependent changeover times and costs along with a more extensive objective function. This formulation was further extended by Kopanos & Puigjaner (2011) to also include manpower constraints and other constraints that are specific to different product families (e.g., packing rates, setup times). The same authors present a model for a multiproduct multistage ice cream production factory (Kopanos & Puigjaner, 2011). This formulation includes strong integer cuts that allow the problem to be solved to optimality for a real-life industrial case. The formulation is further enhanced by introducing new sets of tightening constraints resulting in low computational requirements and solutions of very good quality (Kopanos & Puigjaner, 2012).

The different studies show the recent development of models that are more applicable to real-world and large-scale problems. However, the models also show a tendency to become more specialized towards the process that is being modeled. This makes the presented models less useful for use in the production-scheduling problems that arise in other production processes.

4 Motivation

Research into the problem is motivated from both the operating and the scientific point-of-view. In this section both these motivations will be explained.

4.1 Company point-of-view

For the company, the increasing complexity and the decreasing free capacity of the production process mainly motivates research into this topic. Over the year of 2012 a healthy growth of more than 20% was realized. The same is expected for the year 2013. Although production capacity is in the process of being extended, lead times for new machinery are often long. This puts a strain on the current available capacity and requires more efficient production planning. Furthermore an increase in available machinery increases the number of possible production schedules. Currently, the schedule is only altered when problems are encountered. At that time, a lot of potentially good scheduling options have already been lost. Without an optimization approach, the gap between what a human can manually schedule and the solution of an optimization approach will only grow larger.

As mentioned in the introduction, in 2007, an attempt was made to implement an APS module into the existing ERP system. After a long implementation period it was concluded that the system was not able to cope with the specific needs of the company. Dependency between the different batches and the zero lead-times of a lot of the products could not be accurately modeled. The APS system was unable to incorporate all the constraints of the process and was only able to incorporate some general objective functions. These general objectives did not match the actual objectives the company has for the optimization of the production process. Without the proper way to model the process, the schedules produced by the module were illogical and unfeasible. As a result, the implementation was halted.

Even though the attempts at a more sophisticated approach to the problem did not result in a workable solution, the managers realize that the current manual approach will become insufficient in the near future. This research should provide an insight into the problem and provide a possible solution approach for the problem.

4.2 Scientific point-of-view

The literature review showed that planning and scheduling problems are often complex of nature and very process specific. Although some attempts at general frameworks have been made, these often lack the specialization required for practical real-world applications. It also showed that although some research of similar production processes in the dairy industry exists, no research into fresh vegetable production processes was found.

The fresh vegetable processing industry provides a very interesting challenge because of the short lead-time and the very high fill rates of orders. Any sales that are not fulfilled on time are directly lost. Apart from having to produce every product daily, this means that the production should fulfill every order completely and that production cannot be postponed.

Furthermore, the production process itself provides some interesting and complex challenges. Because of cross-contamination risks, some products can only be produced and packaged in a specific order. If this order is not kept, extensive cleaning of the machine has to be performed. This usually means that less-colored vegetables like white cabbage and cauliflower have to be produced before more colored vegetables like red cabbage and beets. Other challenges of the process include the limited capacity of the buffer inventory and the rules set for the different intermediate products. Because of these rules, all intermediate products have a limited time frame in which they have to be packaged in order to preserve their freshness. Especially these points make the process different from the more conventional production processes often found in literature.

Apart from challenging modeling aspects, the solving of scheduling problems in general is considered to be very hard. Although over the years a lot of different attempts have been made to solve these problems quickly, no polynomial algorithms that provide optimal solutions exist today. This may mean that solving this large scale real-world problem will require a heuristic solution approach or that suboptimal solutions will have to be considered.

This research will attempt address the lack of knowledge of production scheduling in fresh vegetable production processes. Specifically, a better understanding of sequence restricted scheduling and intermediate products with lifetime restrictions should be provided. This will be done by formulating a specialized mathematical formulation of the problem. After this, an attempt should be made to solve the problem using a commercial solver, a fast heuristic or a combination of both.

5 Data

In order to successfully model the productions process, a lot of different data will be necessary. Although some data was easily obtainable from files, a lot of it was constructed with the help of experienced workers at Heemskerk because it was not digitally available. This section will discuss the necessary data for the model, their origin and their accurateness.

5.1 Batches / Orders

The main parts of a scheduling problem are the orders that are required to be scheduled. One of the first decisions to make is which part of the orders should be scheduled in a single instance of the problem. At Heemskerk, production is scheduled on a daily basis. Every day, orders for the production day of tomorrow are prepared. Furthermore, due to the very short lead-time of orders, the entire range of products is produced every day. This effectively means that when we look at a single instance of the problem, every day is a repetition of the day before. It is therefore sufficient to model only the orders of one day in a single instance of the problem. The scheduling horizon of the model will thus be equal to one day.

The short planning horizon should result in a lower number of orders that needs to be scheduled. However, the number of orders is still quite large. Because the entire assortment of products is produced every day there is no reduction in the variety of the orders even though the planning horizon is so small. In table 5.1 we list some of the main characteristics of a typical production day at Heemskerk

Table 5.1: Characteristics of a typical production day at Heemskerk

Description	Amount
Total number of orders	942
Total number of orders for final products	710
Total number of orders for intermediate products	232
Amount of work centers in use	31
Amount of packaging lines in use	16
Amount of pre-processing lines in use	15

Fortunately these numbers can be greatly reduced when a feasible schedule for the production process needs to be created. Every customer has his or her own range of final products. For each of these products, an order is entered into the ERP-system resulting in the large number of orders shown in table 5.1. Often the only difference between these products is the label that is put on the package. The rest of the product is exactly the same. This means that the product, the amount of product per package and the packaging are equal. Whenever the label is removed and another label is put on the package, it can directly be sold to another customer. Since changing the label used on a production line only takes a short amount of time (< 1 minute) compared to the actual production of a product, we can combine all orders that differ only by the kind of label used on the package. This results in a large reduction in the amount of orders that needs to be scheduled.

The number of orders can be further reduced because some of the required intermediate products are produced in the off-hours in advance of their respective production day. These are usually products that have longer lifetimes once they have been processed. Pre-processing lines are often finished before the end of the day. In these off-hours the strong products are produced to stock on the day before they are actually required. This means that on a production day, these products are already available and do not have to be scheduled in the scheduling problem. Therefore the orders for these products can be removed from the problem instance.

Although this practice removes these orders from the problem instance, the amounts required of these specific products on the next production day will be added to the workload. Fortunately, this is only a small amount of the total pre-processing orders and they have very little interchangeability between the different machines. Furthermore, there are no other batches that are dependent on these orders because their respective packaging orders are scheduled for the next day. This means they can be scheduled anywhere in time.

Because these independent orders will only be scheduled after all dependent orders have been finished, this essentially results in a number of small scheduling problems of at most a couple of batches that need to be scheduled in the time available after all other dependent order have been scheduled. These small problems, e.g. the scheduling of two or three batches on a single work center, can easily be solved by hand. This is further supported by the fact that even on one of the busiest days, the total workload of the pre-processing lines contributes to only 48% of the available production time (see section 8.3 for an example). It is therefore assumed that the batches concerning these products can be completely left out of the problem.

Removing some of work centers from the problem instance can also further reduce the number of orders. This will be further explained in section 5.2.

Throughout this thesis a single representative production day will be used to show the results that can be achieved. This day is the 11th of July 2013. This day is a Thursday and is usually the busiest or 2nd busiest day of the week. Therefore results for this day should give a good representation of the possibilities of the model. In table 5.2 the details of this production day are listed. These details include the reduction in the number of orders explained above.

Table 5.2: Production details of the 11th of July 2013 including the reduction in the amount of orders

Description	Amount (difference after reduction)
Total number of orders	357 (-584)
Total number of orders for final products	234 (-476)
Total number of orders for intermediate products	123 (-109)
Amount of work centers in use	27 (-4)
Amount of packaging lines in use	16 (0)
Amount of pre-processing lines in use	11 (-5)

The reduction in the amount of orders for final products is entirely realized by the combination of orders that have the same label while the reduction of orders for intermediate products is mainly realized by leaving out orders on 4 work centers.

To retrieve all orders for the production day of 11-7-2013 the demand of final products was taken from the ERP system. After that, an MRP (material requirements planning) script was used to explode all the requirements from the bill-of-materials into a list of required intermediate products. Together with the orders for final products, these intermediate orders form the list of orders that will need to be scheduled by the model.

5.2 Work centers and available work centers for each product

As was mentioned in section 5.1, orders need to be scheduled on a large number of work centers. There are, in total, 31 work centers that see regular use. Next to that, a number of small orders are performed by hand. These orders can be left out of the scheduling problem because they do not form a bottleneck. Whenever these orders for intermediates are required in the packaging process, they can be processed on time.

A small number of work centers are far from being at their maximum used capacity. These are small washing and mixing machines. Like the orders that are performed by hand, orders assigned to these work centers are no bottleneck for the scheduling process. Plenty of capacity is available here and using extra manpower can increase it further. Therefore, orders on these machines can be left out of the problem instance. We can assume that intermediates coming from these machines are always available in time. Removing the orders on these machines results in the reduction of orders shown in table 5.2.

In table 5.3 we list all the work centers that are included for the scheduling problem.

Table 5.3: Type and availability of scheduled work centers at Heemskerk

Type of work center	Amount available	Descriptions of work centers
Pre-processing line	11	A01, A02, A03, A04, A05, A07, A09, A11, A13, A15, A20, C01
Bag-packaging line	11	B00, B01, B03, B04, B05, B06, B07, B08, B09, B10, B11
Bowl-packaging line	3	B12, B13, B14
High-Care line	2	HC01, HC02

We can distinguish 4 main types of work centers. The first type includes pre-processing machines. These machines are used to turn unprocessed vegetables or other intermediates into intermediate products that are ready for packaging on the 3 different types of packaging lines. These lines have little flexibility. They are equipped to produce a certain type of intermediate. Only in rare occasions these machines are used for product types for which they are not designed.

The first type of packaging line is the bag-packaging line. There are 11 of these work centers available and they are highly flexible. These lines package intermediates into bags and seal them. By exchanging the format part and the packaging film they can produce a wide range of different end products. Orders that are produced on these machines are often changed from their standard work center assignment.

The second type of packaging line is the bowl-packaging line. On these lines intermediate products are dropped into bowls and then sealed by a plastic cover on top of the bowl. This is a more specialized type of packaging than bag packaging and this results in less flexibility between the different lines. The standard machine assignment for orders on these machines is also the only possible machine assignment.

The last type of production line includes 2 High-Care lines. Essentially these lines are equal to the bowl-packaging lines in their form of packaging. However, next to vegetables they also package meat, nuts, pasta and other non-vegetable products alongside the vegetables. Because of cross-contamination risks they are produced in a separate area of the factory where extra hygienic rules are applied. There are 2 of these packaging lines in the factory and they are both identical. They can both produce all products that require a High-Care line.

It is important to note that there is no interchangeability of orders possible between the different types of production lines. If an order is assigned to a pre-processing line it can never be produced on another type of line. However, assigning orders to a different work center may be possible when it is within the same type of machine.

In table 5.4 an overview is given how many of the orders required on 11-7-2013 are assigned to each type of machine. Along with this the amount of these orders each work center can produce is listed on the row for each machine.

Table 5.4: Table listing the amount of orders per type of machine and the possible amount of orders per machine on 11-7-2013

Machine Type	Machine	Amount of (possible) orders
Bag-packaging Line	Total amount of orders:	147
	B00	67
	B01	73
	B03	20
	B04	102
	B05	54
	B06	86
	B07	100
	B08	87
	B09	92
	B10	71
	B11	110
Bowl-packaging Line	Total amount of orders	57
	B12	13
	B13	19
	B14	24
High-Care Packaging Line	Total amount of orders	33
	HC01	33
	HC02	33
Pre-processing line	Total amount of orders	122
	A01	17
	A02	10
	A03	7
	A05	17
	A07	16
	A09	3
	A11	9
	A13	25
	A15	7
	A20	8
	C01	3

As becomes clear from the table, the bag-packaging lines and high-care lines are the most flexible type of line. Along with the high number of orders on the bag-packaging lines, this will results in the highest number of possible scheduling options.

All the data presented in this section was gathered from the companies' ERP system and with the help of an experienced production team leader. He provided all possible alternative machine assignments because the ERP system currently only contains the standard routing for each product. The data was finally translated into a binary matrix holding all products on the rows and all the work centers in the columns. A value of 1 indicates that a certain product/work center combination is allowed. A value of 0 indicates it is not.

It is important to note that the standard assignment in the companies' ERP system can be viewed as the preferred assignment for each product. In case the workload of a work center is not above of what the machine is capable of daily, a product should preferably be assigned to its preferred work center. The difference between the preferred production line and other production lines might not always be that large but it can be assumed that in most cases it is better to assign a product to its preferred machine.

5.3 Changeovers

One of the main disturbers of the production process is a changeover between the different products on the different work centers. In order to more accurately model the process it is necessary to gather the data on these changeovers. Changeovers that are considered to be bad result in extra work for workers on the shop floor. Naturally, introducing a lot of these bad changeovers results in a lot of opposition from the workforce. Apart from this, having a production schedule with a lot of difficult changeovers results in longer processing times.

The changeover process on pre-processing lines is very different from the process on packaging lines. Therefore, they will be listed separately here.

5.3.1 Changeovers on pre-processing lines

Whenever the order for an intermediate product finishes, the production line may have to be cleaned. On this type of production line, the amount of cleaning required to produce the next product according to its specifications indicates how bad the changeover actually is. For the pre-processing lines three types of changeovers can be identified. These three types are listed in table 5.5.

Table 5.5: Types of changeovers on pre-processing lines

Description of changeover type	*Amount of time (min) required for changeover*
1. Emptying machine and/or changing the type of cut made by the machine	5 minutes
2. Light cleaning of the machine	10 minutes
3. Thorough cleaning of the machine	45 minutes

In order to indicate which type of changeover occurs when the work center changes over to a different product, a large matrix was created that lists the type of changeover that occurs for each possible combination of two products. This means that when a type 2 changeover occurs when changing from product *i* to product *j,* the value in the matrix in cell (*i,j*) equals 2.

It is important to note that if a type 2 changeover occurs when switching from product *i* to product *j,* this will usually not be the same when switching from product *j* to *i*. In many cases when it is easy to switch from *i* to *j,* it is a lot more work to switch from *j* to *i*. This means that the changeover matrix is not symmetrical. As a result of this, the choice of event representation in the model will be limited (Baumann & Trautmann, 2013). This will be further explained in section 6.1.3.
Figure 5.1 shows a small part of a changeover matrix for pre-processing line A01.

		59005 Kastanje champ rb mengsel A01	50131 Chinese gr mix gew gedr	50260 Chinese rb mengsel	50353 Mac spag meng 8mm A01	50196 Bamigroenten nasi mix 6mm	50394 Bamigroenten mengsel 6mm SU	50090 Bamigroenten mengsel	50223 Hong Kong rb mengsel A01	50130 Italiaanse rb gew gedr
59005	Kastanje champ rb mengsel A01		2	2	2	2	2	2	2	2
50131	Chinese gr mix gew gedr	3		1	2	2	2	2	2	2
50260	Chinese rb mengsel	3	3		2	2	2	2	2	2
50353	Mac spag meng 8mm A01	3	3	3		2	2	2	2	2
50196	Bamigroenten nasi mix 6mm	3	3	3	3		1	1	2	2
50394	Bamigroenten mengsel 6mm SU	3	3	3	3	1		1	2	2
50090	Bamigroenten mengsel	3	3	3	3	1	1		2	2
50223	Hong Kong rb mengsel A01	3	3	3	3	3	3	3		2
50130	Italiaanse rb gew gedr	3	3	3	3	3	3	3	3	

Figure 5.1: Part of the changeover matrix for preprocessing line A01.

5.3.2 Changeovers on packaging lines

Changeovers on packaging lines mainly depend on the recipes of the products. Three components of the bill-of-materials of a product determine which type of changeover occurs when switching between products. These 3 components are the film used for packaging, the kind of intermediate vegetable used, and the label on the packaging.

Whenever the film of a product is different from the previous product, the roll of film has to be changed. The type of film also directly relates the format part that is used on a machine. The format part determines the size of the package. A package that contains 600 gr of product requires a larger format part than a package containing only 200 gr of product. A larger format part requires a wider roll of film. Therefore, the type of film used in the bill-of-material of a product can be used to determine if a format part exchange is required.

If the intermediate product used in the recipe changes significantly when the product on a machine is changed, the machine should be cleaned. This means that leftover parts of the previous product are removed with compressed air.

Last, when the label on a product is different from the product previously produced, the roll of labels in the label printer has to be changed. Orders are set up in such a way, that a different label is the smallest change that results in a different order. This means that on every order change the type of label has to be exchanged. Therefore the time required for this can be included in the processing time. Due to this, as was mentioned in section 5.1, we can combine all orders that differ only by the type label used into a single order. This significantly reduces the number of orders that are required to be scheduled.

The required changes described above result in three distinguishable types of changeovers. Whenever the type of product changes, the machine has to be cleaned. When the type of film changes, the roll of film has to be exchanged and when the size of the film also changes, the format part of the machine has to be exchanged as well. In table 5.6 we list these operations and the times required to finish them.

Table 5.6: Type of changeovers and their respective durations.

Type of Changeover	Duration
Change of format part	10 minutes
Change of film	5 minutes
Change of product	4 minutes

The change of product is an independent event. However, the change of film and format part are related. When a format part changes, the film always has to be changed as well. When the type of film is changed, the format part does not necessarily have to change as well. This results in a limited number of changeover combinations. These combinations are listed in table 5.7. An x denotes that this type of changeover occurs in this combination.

Table 5.7: Possible changeover combinations on packaging lines

Changeover #	Format change	Film Change	Product change
1	x	x	x
2	x	x	
3		x	x
4			x
5		x	

The changeover matrix, which denotes what changeover occurs when product *j* is produced after product *i,* was created using a script that compares the bill-of-materials of all products to each other. In this comparison the changeover number (table 5.7) is retrieved for each possible combination of two products.

5.4 Processing and other durations

To correctly schedule the orders, the processing time required to finish each order is required. Unfortunately the processing times at Heemskerk are not frequently updated and stored in a single database. However, there is extensive transaction data available. A transaction is a data entry that includes: the time and date at which the transaction occurred, the amount and type of product that was designated as completed by the transaction, and the work center on which the transaction occurred. With this data it is possible to calculate how much time was used to complete a certain amount of a product.

This transaction data was gathered over a period of three months. This resulted in a total of over 250.000 transactions. This data was further refined by removing workforce lunch breaks and errors. From this data estimated processing times required per unit were calculated. These processing times were made dependent of the size of the orders. A larger order results in shorter processing times per unit. These processing times also include the time required for changeovers.

Since then, these processing times have been used in expected workload overviews and capacity planning software. The overall view of the workforce at Heemskerk is that these times quite accurately describe the actual time required to produce the orders. They are now also constantly being updated with new transaction data.

Another important time frame is the time required before an intermediate product can be used in another product. This 'quarantine time' will be included as a restriction in the model. This time is required to transport the products from their respective production line through the buffer inventory to the production line where they are required. Production of intermediates happens at a greater rate than the rate at which they can be used up. Furthermore, intermediates are stored into small metal containers on wheels. This way orders for intermediates that are only partially finished can already be used in the next step of the production process. Therefore, it will be sufficient to have started production of any intermediates at least 30 minutes before they are required.

5.5 Contamination levels

An important aspect of the production process is the prevention of cross-contamination. A product is defined as contaminated when the product contains visual traces from other products or does not taste the way it should. To prevent this there are clear restrictions on the order in which products are allowed to be produced. Every night, after the production day has finished, all production lines are thoroughly cleaned (>1 hour). This means that at the start of a new production day, all packaging lines are clean and everything can be produced. As soon as a product with a certain contamination is produced on a packaging line, a range of products can no longer be produced without being contaminated. These products are only allowed to be produced again after the production line has been thoroughly cleaned during the next night.

A clear example of this; is the packaging of white onion rings. Any product that is not white as well easily contaminates this product. If a green-colored product like spinach is produced before this product, green parts will easily contaminate the packaging of the onions and they will be clearly visible. It is clear that these white onion rings should always be produced before anything that is not white.

In order to correctly model these restrictions, a list containing the contamination level for each product was composed. In this list, products with a high contaminations risk are assigned a high level while products with a lower contamination level are assigned a lower contamination levels. Once a product with a certain contamination level is produced, only products that have an equal or higher contamination level are allowed to be produced.

The onion rings discussed in the example have a contamination level equal to 1. The spinach has a contamination level of 3. When enforcing the rule above, the spinach can only be produced after the onions. Figure 5.2 shows an example of the contamination levels on the B04 packaging line.

Item Nr	Description	Standard Line	Contamination level
11844	Uien 1000gr	B04	1
10536	Bloemkool 350g	B04	2
10538	Broccoli 300g	B04	2
11535	Andijvie fijn 200g	B04	3
11215	Andijvie 250g	B04	3
11562	Andijvie fijn 400g VV	B04	3
13594	Andijvie fijn 500g VV	B04	3
11867	Andijvie fijn 450g VV	B04	3
11120	Andijvie 600g VV	B04	3
11211	Andijvie 400g	B04	3
11275	Saladinettes Andijvie 400g	B04	3
11256	Babyleaves 200g	B04	3

Figure 5.2: Contamination levels of products with a standard assignment to the B04 packaging line.

The contamination levels were determined in a dialogue with a production team leader. On each production line, all products that can be produced on that line were taken into account.

The contaminations levels are only used on packaging lines because these cannot be cleaned sufficiently during a production day. On pre-processing machines this is not a problem since cleaning takes less time and can be performed between production runs.

5.6 Available manufacturing timeframe

The process at Heemskerk knows a fixed preferred ending time at which all production should be finished. This allows for the overnight cleaning of the production lines to take place. Only in rare occasions where the workload is simply too high to be finished within this timeframe, production lines are allowed to produce after this ending time.

Production happens during two shifts. Each shift includes 9,5 hours of production time and 1 hour for breaks. This results in a 4-day workweek of 38 hours for the production personnel. The total production time available on each work center is thus equal to 19 hours per day. With rotating shifts, the production facility is opened for all 7 days in a week.

Pre-processing production lines start their first shift at 5:15am while packaging machines start their first shift at 6:15am. Due to this time shift, pre-processing lines are allowed to build up a small buffer inventory of intermediates to be used for the first production runs on the packaging lines.

The second and last shift of the pre-processing and packaging lines ends at 1:15am and 2:15am respectively. After this, the thorough cleaning of the production lines can commence.

6 A mathematical optimization model for the scheduling problem

In this section, a mathematical optimization model for the problem will be proposed. The choice of model type will be defended and the incorporation of the key process constraints in the model will be explained. It will also be shown how the data presented in section 5 will be included in the model description.

6.1 Choice of model type

As was mentioned in the literature review, production-scheduling problems appear in a wide range of different types. Each type of model has its own strengths and weaknesses but some are more suited for certain processes than others. Here, the decision is made which type of model to use for the problem. Several independent choices lead to the final decision.

6.1.1 Discrete- or Continuous-time representations

The first choice for every scheduling model is whether time should be represented on a discrete or on a continuous time-scale. A discrete time-scale will result in a simpler model but the time representation will only be as accurate as the smallest chosen time interval. For this type of model, the number of constraints and variables will increase when the size of the time intervals becomes smaller. On the other hand a continuous time scale will have an accurate representation of time but the model may be more complicated and require more sophisticated constraints.

At Heemskerk the shortest time required per order is 2 to 5 minutes. A typical production day should finish within a timeframe of 21 hours. However to leave a large enough feasible region (to allow to scheduling in overtime), a timeframe of at least 36 hours would be required. Even taking a relatively inaccurate timeframe of 10 minutes would already lead to 216 intervals. Stefansson & Sigmarsdottir (2011) furthermore show that scheduling a discrete model that has more than 100 orders already results in solution times greater than 60 minutes. The continuous model they present performs significantly better than the discrete model for any model size of over 100 orders.

The time intervals required to accurately model the process at Heemskerk result in a large amount of intervals for the discrete model. Furthermore, the amount of orders that needs to be scheduled far exceeds 100 orders. This makes the choice for a continuous-time representation easy to defend.

6.1.2 Material Balances

According to (Mendez, Cerda, & Grossmann, 2006) the handling of batches and batch sizes gives rise to two types of optimization model categories. The first type refers to approaches that simultaneously deal with an optimal set of batches (number and size), the allocation and sequencing of manufacturing resources and the timing of processing tasks. Including all of these into one model usually results in large model. This restricts their use to processes involving only a limited number of tasks. These models use network representations to model the flow of materials. The two types of network representations are the state-task network (STN) and the resource-task network (RTN).

The second category includes models that assume that the number of batches of each size is known in advance. This is a type that allows wide use in industrial applications. First the batching problem is solved separately. After this the scheduling problem is solved where the batches are assigned to the available manufacturing resources. This two-stage approximate approach is able to cope with much larger problems than the first model type.

At Heemskerk the production is batch-oriented. Since each day almost every end product in the assortment is produced, the batch sizes are not very large and do not often need splitting. Since the planning horizon is only one day, producing the same product multiple times within that same day is not very efficient. Introducing the additional changeovers required does not outweigh the benefits of producing products twice or more per day. Next to that, the restrictions in place to prevent cross-contamination often do not allow products to be produced again at a later time.

The process at Heemskerk is therefore well suited for the employment of a batch-oriented model from the second category. Because orders are not split, the batch sizes are known in advance. Furthermore the amount of orders that need to be scheduled will most likely be too high for any network-based representation to handle.

6.1.3 Event representation

The batch-oriented continuous-time model that will be used for the problem allows events to be represented in one of three different ways. These event representations mainly differ in the way in which the allocation and sequencing variables are defined.

In the first type of event representation, unit-specific immediate precedence, both the sequence of the production and the machine assignment are represented by a single variable. The binary variable $X_{ii'j}$ then denotes if order *i* is processed right before order *i'* on work center *j.* This also requires an additional variable in the form of XF_{ij} which denotes if order *i* is the first processed order on machine *j.* This type of event representation is very specific and allows for more specific constraints and objective functions. It does however result in a large amount of variables since sequencing variables exist for every separate work center.

The second type of event representation, immediate precedence, separates the variables for machine assignment and sequencing. Here, $X_{ii'}$ denotes if order *i* is processed right before order *i'*. F_{ij} / X_{ij} denote if order *i* is the first / is assigned to work center *j.* Splitting these variables results in a large reduction of variables without losing much details.

In the final type of event representation, general precedence, $X_{ii'}$ denotes if order *i* is processed before or after order *i'*.In this type of representation, the immediate relationship between orders that are processed right after each other is lost. However the variable F_{ij} is no longer required in this representation. This simply halves the number of assignment variables required compared to the immediate precedence representation. If no additional constraints have to be added to accompany this change, this representation will thus result in the lowest solution times.

In the model for the production process bad changeovers will be penalized. In order for this to work with a general precedence event representation, the changeover matrix should satisfy the triangle inequality (Baumann & Trautmann, 2013). This means that the penalty for switching over from product *i* to product *j* should always be lower than the penalty for first switching over from product *i* to product *j* via product *k.* At Heemskerk this is not the case as a result of the cross-contamination restrictions. The general precedence event representation can therefore not be used as an accurate representation.

In contrast to general precedence, an immediate precedence event representation will allow for accurate modeling of the exact changeovers between products. Therefore it will be used for the model. A unit-specific immediate representation would increase the amount of variables without giving a clear advantage over immediate precedence.

6.2 Including the company specific constraints

A number of process-specific constraints that limit the production process should be incorporated into the model. Some additional constraints will be required to model an objective function that is in line with the objectives of the company. How this was done will be explained in this section.

6.2.1 Modeling changeovers

One of the main objectives of production scheduling at Heemskerk is the minimization of bad changeovers. The main aspect that qualifies a changeover as bad is the willingness of an operator to perform that changeover. Exchanging the roll of film or the format part on a production line is heavy work and no easy task. Introducing a lot of these changeovers will greatly lower the morale of the operator. When a format part only has to be exchanged once or twice a day, the operator will feel that he is producing very efficiently and will try to finish the changeover as quickly as possible. On the other hand, when a format part has to be exchanged 4 consecutive times in a row, the operator may no longer feel the necessity to hurry up because he is not producing efficiently anyway.

Another aspect of changeovers is the time they require. On paper, they may not require a lot of time but when the amount of them is increased they can sum up to a considerable amount of time.

In order to model changeovers, changeover penalties are introduced to the model. Good changeovers will have a lower penalty than bad changeovers. Using a penalty will make it easier to implement this objective into the model. It will also allow for more steering. If a certain changeover turns up too many times in the model results, the penalty can be increased. The drawback of this is that we cannot model the changeover durations directly into the required processing times. The influence of this choice will however be limited. The actual time saved by reducing bad changeovers does not contribute to a large part of the total required processing time.

For example, the B00 production line has an average of 7 product exchanges, 7 film roll exchanges and 2 format part exchanges per day. According to table 5.6 these will results in a total time required for changeovers of 86 minutes. Part of these changeovers will be mandatory to be able to produce all products assigned to that line (3 different format parts per day would already require at least 2 exchanges). Therefore the model can influence only a small part of the total changeover time.

Furthermore, as will be shown later, one of the main objectives of the scheduling problem is for each work center to be finished within a certain time frame. Since the used processing times will include average changeover times required in the current situation, and the model will be equal to or improve upon this current solution, this time frame will never be exceeded by the solution that the model provides.

Therefore using penalties for changeover instead of changeover times should better represent the objectives of the company without introducing extra inaccuracies.

6.2.2 Modeling cross-contamination restrictions

In section 5.5 the cross-contamination rules that apply to the production process were discussed. These rules only apply to batches that need to be assigned to packaging lines.

It was shown that contamination levels were gathered for each packaging batch. These contamination levels were set up in such a way that they are valid for every possible packaging line to which a certain product can be assigned. This means that if onions are assigned contamination level 1 for the B00 production line, they will also have a contamination level of 1 for the B10 production line.

The rule that should be incorporated in the model to prevent cross-contamination is stated in (6.1).

On work center j an order i' with a contamination level equal to k is only allowed to start after order i when $k \geq l$ *. where l is the contamination level or order i* **(6.1)**

In the first few iterations of the model, the cross-contamination restrictions were included in the form of big M-constraints. This constraint prohibited any order *i* that is assigned to work center *j* and which has a lower contamination level than order *i'* from starting after the finishing time of order *i'*. However, big M-constraints are an inefficient modeling practice often resulting in bad LP-relaxations.

The final version of the model will include rule (6.1) by fixing the sequencing variables to 0. This means that all sequencing (binary) variables that result in orders with a lower contamination level being sequenced after something with a higher contamination level are fixed to 0. These variables effectively become parameters and therefore result in a much more efficient model.

The method of fixing the sequencing variables results in a large reduction in the number of variables and constraints. The difference this makes for the model instance of the exemplary production day of 11 July 2013 are shown in table 6.1

Table 6.1: Difference between using big-M constraints and fixing the variables to model cross-contamination restrictions.

Method used	# of Variables	# of Constraints
Big M-constraint	139,266	650,384
Fixing sequencing variables	17,944	529,677
% difference	**-87.11%**	**-18.56%**

6.2.3 Modeling time restrictions

In section 5.6 the timeframe during which production is allowed to be scheduled is explained. Introducing this time limit as a hard constraint limits the available scheduling interval for the model. Any solution that does not finish exactly within this timeframe would be infeasible. Furthermore, solutions that fall just outside the time interval but may still be acceptable would be infeasible as well. Also it would not be possible to get a possible solution for problem instances that have no feasible solution that fits within the timeframe.

Because solutions that do not fit exactly into the restricted timeframe should still be shown and because these solutions may be improved iteratively later, the production time frame restrictions will be introduced to the model as soft constraints.

Orders that are scheduled by the model to finish after the allowed production timeframe will be penalized in the objective function. The weight of this part of the objective function can then be varied to modify the degree in which the constraint is allowed to be violated.

Generally, only solutions that (almost) fit into the production timeframe should be accepted as solutions to the problem.

6.2.4 Modeling the dependencies between intermediates and end products.

Before final products can be produced, all required intermediates should be available. In section 5.4 it was mentioned that production of orders that are required for other orders should start at least 30 minutes before the starting time of said order.

Creating subsets of orders that are required for each order enforces this constraint. The constraints are then enforced only for the orders that are in this subset.

6.3 A mathematical formulations for the production-scheduling problem

The presented data and process restrictions can be used to formulate a mathematical optimization model for the scheduling problem. Here, an immediate precedence, continuous-time batch scheduling model is presented. The model can be used to simultaneously schedule both the pre-processing and packaging orders at once.

In section 6.3.1 on this page the mathematical formulation will be given. In section 6.3.2 the sets, variables and parameters will be explained, in section 6.3.3 the constraints will be explained and in section 6.3.4 the objective function will be discussed.

6.3.1 Mathematical formulation

Sets

W	Batches / Work Orders
$W^m \subset W$	Make batches (batches that produce intermediates)
$W^p \subset W$	Pack batches (batches that produce final products)
$W_w^m \subseteq W^m$	Make batches that supply pack batch $w \in W^p$
$W_w^r \subseteq W^m$	Make batches required to be finished before make batch $w \in W^m$
J	Work centers (production units)
$J_w \subseteq J$	Work centers that can produce batch $w \in W$
$J_{ww'} \subseteq J$	Work centers that can produce both $w \in W \textit{ and } w' \in W$

Parameters

α_w	Processing time required per unit of batch $w \in W$
δ_w	Time required before batch $w \in W$ can be used in another batch
$\omega_{ww'}$	Penalty for a switchover between batch $w \in W$ and batch $w' \in W$
β_w	Number of units required of batch $w \in W$
M	A sufficiently large number
D_w	Preferred time at which work order *w* should be completed
$K_{ww'}$	Binary matrix that is equal to 1 when for a pair of batches *w* and *w'* the variable $Y_{ww'}$ should be equal to 0.

Continuous Variables

S_w	Start time of batch $w \in W$
D_w^+	Amount of time that batch $w \in W^p$ finishes after the preferred completion time D, i.e., the tardiness of order *w*

Binary Variables

F_{wj}	1, if batch $w \in W$ starts the processing sequence on workcenter $j \in J_w$
X_{wj}	1, if batch $w \in W$ is assigned to but not the first to be processed on workcenter $j \in J$
$Y_{ww'}$	1, if batch $w \in W$ is processed right before batch $w' \in W$

Objective function

$$\boldsymbol{min\ \ p_1 \sum_{w\in W}\sum_{w'\in W} \omega_{ww'} \cdot Y_{ww'} + p_2 \sum_{w\in W} S_w + p_3 \sum_{w\in W} D_w^+} \tag{6.2}$$

Allocation constraints

$$\textstyle\sum_{w\in W} F_{wj} \leq 1 \quad \forall\ j \in J \tag{6.3}$$
$$\textstyle\sum_{j\in J} F_{wj} + \sum_{j\in J} X_{wj} \ \ = 1 \quad \forall\ w \in W \tag{6.4}$$

Sequencing-allocation matching constraints

$$F_{wj} + X_{wj} \leq X_{w'j} - Y_{ww'} + 1 \quad \forall(w \in W\ ;\ w' \in W\ ;\ j \in J_{ww'}) \tag{6.5}$$
$$Y_{ww'} + Y_{w'w} \leq 1\ \forall\ (w \in W\ ; w \in W) \tag{6.6}$$
$$F_{wj} + X_{wj} \leq 1 - Y_{ww'}\ \forall\ (w \in W\ ; w' \in W; j \in (J_w \backslash J_{ww'})) \tag{6.7}$$

Sequencing constraints

$$\textstyle\sum_{j\in J_w} F_{wj} + \sum_{w'} Y_{w'w} = 1 \quad \forall\ w \in W \tag{6.8}$$
$$\textstyle\sum_{w'\in W} Y_{ww'} \leq 1 \quad \forall\ w \in W \tag{6.9}$$
$$Y_{ww'} = 0 \quad \forall(w \in W\ ; w' \in W)\ |\ K_{ww'} = 1 \tag{6.10}$$

Timing constraints

$$S_{w'} \geq S_w +\ \beta_w\alpha_w + \delta_w \quad \forall \quad (w' \in W^m\ ; w \in W_{w'}^r) \tag{6.11}$$
$$S_{w'} \geq S_w + \delta_{w'}\ \forall\ \ (w' \in W^p\ ; w \in\ W_{w'}^m) \tag{6.12}$$
$$S_{w'} \geq S_w + \beta_w\alpha_w - M(1 - Y_{ww'})\ \forall\ (w \in W\ ; w' \in W \tag{6.13}$$
$$S_w \geq 1 \quad \forall\ (w \in W^p) \tag{6.14}$$

Additional constraints

$$S_w + \alpha_w\beta_w \leq D_w + D_w^+ \quad \forall\ w \in W^p \tag{6.15}$$
$$F_{wj} + X_{wj} = 0\ \forall\ (w \in W^p\ ; j \notin J_w) \tag{6.16}$$
$$Y_{ww'} = 0\ \forall\ (w \in W; w' \in W; j \notin J_{ww'}) \tag{6.17}$$

Variable constraints

$$S_w \geq 0 \quad \forall\ w \in W \tag{6.18}$$
$$D_w^+ \geq 0 \quad \forall\ w \in W \tag{6.19}$$
$$X_{wj} \in \{0,1\}\ \ \forall\ (w \in W\ ; j \in J_w\ ) \tag{6.20}$$
$$Y_{ww'} \in \{0,1\}\ \forall\ (w \in W\ ; w' \in W) \tag{6.21}$$
$$F_{wj} \in \{0,1\}\ \forall\ (w \in W\ ;\ \ j \in J_w) \tag{6.22}$$

6.3.2 Sets, parameters and variables.

Sets

The set W contains all batches or orders that need to be scheduled by the model whereas the subsets W^m and W^p contain only the pre-processing (make) and packaging (pack) batches respectively.

A subset W_w^m of the set of make batches W^m is created for every order w. It contains all make batches that are required for order w. The subset W_w^r of the set of make batches contains all batches required in pack batch w.

All work centers are included in the set of work centers J. To indicate which work centers are available for each order, the subset J_w is introduced for each order. It contains the available work centers for order w. A second subset $J_{ww'}$ is the subset of work centers that can produce both order w and w'.

Parameters

The processing time required for each unit of order w is given by the parameter α_w while the amount required of each order w is given by β_w. The total amount of time required for each order will thus be given by $\alpha_w \cdot \beta_w$

The parameter δ_w contains the time before an order w can be used in another order. This is set equal to 0,5 hours for all orders that are processed.

The penalty matrix discussed in section 6.2.1 is contained in the parameter $\omega_{ww'}$ where for each changeover from order w to w' the penalty is given.

The parameter M equals a single sufficiently large number. To test the influence of this parameter, it was varied from 2 planning horizons (48 hours) to 10^9. Any number higher than about 10^6 would result in rounding errors. No significant differences in the final results were found for any of the values that did not give errors. Setting this parameter relatively high (in the order of multiple planning horizons) did however result in a slight reduction in the time required to find a feasible solution (< 10 seconds).

D_w gives the time, in hours, at which order w should preferably be finished. For pre-processing orders the value of this parameter will equal 19 (the shift start at 0 and has a duration of 19 hours) and for packaging orders this parameter will equal 20 (since the 19 hour shift starts 1 hour later).

In section 6.2.2 the inclusion of the cross-contamination rules in the model was discussed. To enforce these rules, certain changeovers are prohibited. The parameter $K_{ww'}$ is equal to 1 if a switchover from order w to w' is prohibited.

Variables

The continuous variable indicating the start time of each batch w is given by S_w while the amount of time that each order w exceeded the preferred finishing time D_w, i.e., the tardiness, is represented by the variable D_w^+.

Machine assignment for order w to work center j happens through the binary variables F_{wj} and X_{wj} Here F_{wj} indicates if a work order is the first to be processed on work center *j* and X_{wj} indicates if batch w is assigned to work center j but not the first to be processed on that workcenter.

Finally, the binary variable $Y_{ww'}$ indicates that order w' is processed right after order w.

6.3.3 Constraints

Allocation constraints

Constraints (6.3) ensure that only one order can be the first processed order on each work center. To make sure that all orders are assigned to a work center, constraints (6.4) are introduced. For these constraints an order should either be the first (F_{wj} = 1) assigned or not ($X_{wj} = 1$).

Sequencing-allocation matching constraints

To ensure that all orders allocated to a work center are also assigned valid production sequences, a number of sequencing-allocation matching constraints are introduced. Constraints (6.5) ensure that when two batches w and w' are linked through the immediate precedence constraints $Y_{ww'}$ they are also assigned to the same work center. Constraints (6.6) enforce that a batch w cannot be both the successor and the predecessor of batch w'. Constraints (6.7) then ensure that batches that are not linked through the immediate precedence variable are not assigned to the same work center.

Sequencing constraints

The following constraints ensure that the sequence of all orders is valid and makes sense. Constraints (6.8) make sure that a batch w is either the first order to be processed on a machine or that it is the successor of another order. Constraints (6.9) ensure that the order w can at most be succeeded by one other order w'. Finally, constraints (6.10) make sure that all sequencing variables $Y_{ww'}$ are equal to 0 for pairs of orders for which $K_{ww'}$ is equal to 1. These last constraints are used to enforce the contamination restrictions of the process that were discussed in section 5.5.

Timing constraints

The timing constraints are used to enforce the correct timing of the orders on the continuous timescale used. Constraints (6.11) ensure that make batch $w' \in W^m$ only starts after all its required components $w \in W_{w'}^r$ have been finished and available for at least δ_w hours.
Constraints (6.12) then enforce that pack batch $w' \in W^p$ starts at least δ_w hours after its required make batches $w \in W_{w'}^m$ have started production.

The third type of timing constraints (6.13) ensures that batches that are assigned to the same work center cannot overlap in time. When batch w' is linked to batch w through the immediate precedence variable $Y_{ww'}$ then batch w' is only allowed to start after the finishing time of batch w.

The last type of timing constraints (6.14) enforce that pack batches start one hour after the lowest possible starting time of make batches (=0). This is the result of the one hour difference between the starting time of the pre-processing shift and the packaging shift.

Additional constraints

To complete the entire model some additional constraints are required. Constraints (6.15) are used to calculate the D_w^+ variable. Since D_w is a fixed parameter, variable D_w^+ has to be raised when order w finishes after the desired finishing time.

Constraints (6.16) are tightening constraints that ensure that the assignment variables X_{wj} and F_{wj} are equal to zero for orders *w* that cannot be processed on work center *j* (work centers that are not in the set J_w).

Constraints (6.17) are additional tightening constraints that ensure the immediate precedence variable is equal to zero for combinations of orders w and w' that have no common work center on which they can be produced. In other words, there are no work centers j in the set $J_{ww'}$.

Constraints (6.18) to (6.22) are standard variable constraints

6.3.4 The objective function

Production scheduling at Heemskerk is subject to a number of main objectives:

- *The connection between pre-processing and packaging should be fluent. This means that whenever packaging of a product is scheduled to start, its components should be readily available.*
- *Production should be fluent. This means limiting the amount of changeover to only the ones that are absolutely necessary*.
- *The amount of idle time on packaging lines should be as low as possible. Operators on packaging lines are specialized workers that should not have to fill their shifts with work other than operating their production line.*
- *All production lines should have finished their work before the end of the last shift (before the preferred finishing time D).*

The first objective is included as a hard constraint in the model. Constraints (6.11) and (6.12) ensure intermediates are readily available when the orders consuming them start.

The objective function of the model is a weighted sum. It consists of 3 separate parts which each have their own separate weight factor (p_1, p_2 and p_3). These weight factors can be used to steer the results of the model solution. Giving one of the factors a higher weight than the other should result in a more apparent result of that optimization aspect in the solution.

The first part of the objective function, the part associated with weight p_1, limits the total changeover penalty. Whenever a changeover from w to w' is scheduled, the variable $Y_{ww'}$ will become equal to 1. This means that a changeover penalty equal to $\omega_{ww'}$ is introduced. Because bad changeovers have higher penalties, the total sum associated with this part of the objective function will become higher when a lot of bad changeovers are introduced.

The second part of the objective function is associated with weight factor p_2. This part is equal to the sum of all starting times. It is used to minimize the amount of idle time that will occur on each line. Without this constraint there is little incentive for the model to schedule all orders right after each other when the total workload does not exceed the preferred finishing time. A clear example of this mechanic is given in figure 6.1. Both bars in this figure represent the same production line with an equal number of the same orders. In both schedules the part of the objective function associated with p_3 is equal to zero. However, there is a lot of idle time (white spaces) in the second schedule. By introducing the second part of the objective function ($p_2 > 0$) this idle time can be effectively limited. The first bar in figure 6.1 shows the result of this.

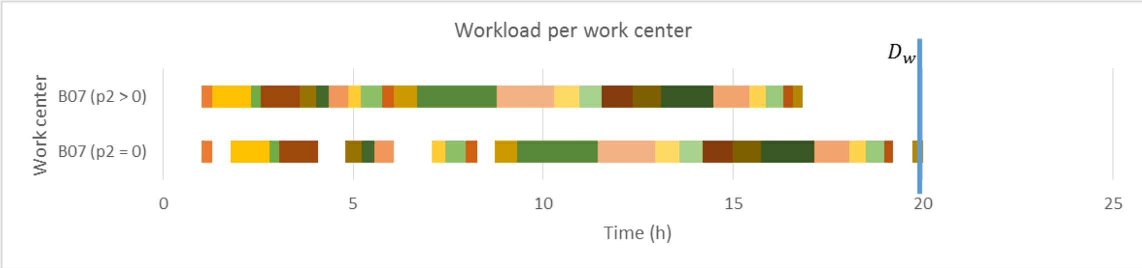

Figure 6.1: Example of a schedule with and without the inclusion of the 2nd part of the objective function.

The total weighted sum of the 3 parts of the objective function includes all the objectives that a good solution to the scheduling problem should have. Since all parts of the objective function are penalties that should be kept as low as possible, the objective function is to be minimized.

The third part of the objective function is associated with penalty p_3. This part is used to minimize the amount of time each order is finished after its preferred completion time. Since D_w^+ is only greater than zero when an order is finished late, this penalty will be equal to zero when no orders finish late. This part of the objective function is used to enforce the last of the objectives at Heemskerk. As was mentioned before in section 6.2.3, this objective is not introduced as a hard constraint. It is introduced into the objective function. Generally only solutions where this part of the objective function is at or very close to zero will be accepted.

Figure 6.2 shows an example where this part of the objective function is not equal to zero. The workload of two work centers is shown. Production line B04 finishes before the preferred delivery time D_w while line B05 finishes after D_w. Therefore, in order to satisfy constraints (6.15), D_w^+ is not equal to zero.

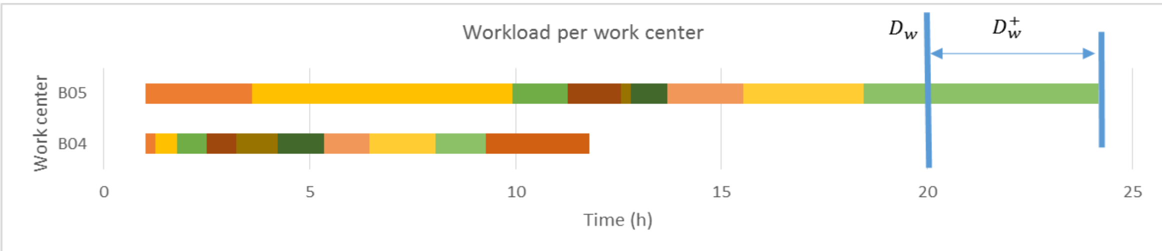

Figure 6.2: Example of a schedule in which one production line exceeds the preferred delivery time.

6.4 Preliminary model results

The first idea was to solve the model presented here for a single instance (one production day) of the problem. This was attempted using the *CPlex 12.5* solver within the modeling software solution *AIMMS*. The size of the problem was gradually increased by introducing more orders into the problem. The first results are listed in table 6.2.

These results were obtained on a notebook with an Intel i5-3317U CPU and 8GB of DDR3 ram.

Table 6.2: List of preliminary results to test the performance of the CPLEX solver at solving the presented model.

# of orders	*Solution Feasible(s)*	*Sol. time (s)*	*# Variables*	*# Constraints*	*Gap %*
22	0.02s	1.83s	1028 (961 integer)	1980	optimal
48	0.53s	1111.23s	3045 (2900 integer)	8661	11.43%
86	58.47s	1667.02s	7838 (7579 integer)	32573	30.39%
357 (all)	n/a	47,160.00s	17944 (16872 integer)	529677	n/a

The solutions show that it is hard to find an optimal solution. The LP bound of the problem is not very tight. This results in greater optimality gaps for problem instances involving a higher number of orders. However, as will be shown in section 8, the feasible solutions the solver obtains are very good. The solutions on line 2 and 3 of table 6.2 were found by letting the solver run until no better feasible solution was found for 600 seconds. For the last problem instance, which involves all 357 orders, no feasible solution was found. After 47,160 seconds the solver ran out of memory to use.

These preliminary results show that it is unlikely that just using a commercial solver on the model described in this section can solve the problem. A feasible solution was not found before the solver ran out of memory. To find an acceptable solution within a reasonable amount of time (< 1-2 hours), a different approach will need to be pursued. This approach will be presented in section 7.

From here on, the approach presented in this section, i.e. simply solving the model presented here using a commercial solver, will be referred to as the 'MIP-based approach'.

7 Solving the model for a large-scale real world problem

Section 6 showed that the usage of a commercial solver to simply solve the model will not be sufficient to obtain acceptable solutions within a reasonable amount of time. In this section a MIP-based sequential solution strategy is introduced that reduces the size of the problem instances. Furthermore, an initial solution is provided which is then iteratively improved using a commercial solver.

7.1 Process aspects that can be used to solve the model

The production process has a few distinct aspects. Knowledge of these aspects can be effectively used to obtain better solutions in a more reasonable amount of time. In this section a number of these key aspects will be listed. Knowledge of these aspects is then used in combination with a commercial solver to find good solutions

7.1.1 Separation between different production departments

At Heemskerk there is a rather strict separation between the three different packaging departments and the pre-processing department. Because the four departments each have their own specialization, and the production lines have been adjusted for this, order assignment cannot be easily switched between the different departments.

A bag-packaging production line is not able to package products into bowls and a bowl-packaging line cannot seal products into bags. The third type of production line, the high care line, is technically able to produce products that are normally produced on a bowl-packaging line. However, it is not possible to produce high care products on the bowl-packaging lines due to the contamination restrictions that apply to the former. Furthermore, because high care products contain more ingredients and therefore require more workers at the production line, it is highly inefficient to produce products from the bowl-packaging department on a high care production line. So although technically possible, this is never done in practice.

Because orders assigned to the three different packaging departments do not depend or interfere with each other, the problem instances of these departments can be split into three instances that are much smaller. This can be done without influencing the solutions for the departments because orders that belong to one department cannot be assigned to work centers in other departments.

The challenge in splitting the problem for the packaging lines into three separate problem instances lies with the pre-processing department. Although production orders for intermediates can only be assigned to pre-processing work centers, the production of intermediates is linked to all other departments. A single intermediate may be needed for a separate production order in all three packaging departments.

Fortunately constraints (6.12) are the only constraints that connect orders that are assigned to the pre-processing department to the other departments. These constraints ensure that the packaging process can only start when the required intermediates are available.

To be able to split the model into 4 smaller problem instances for each production department, two versions of the model will be created. The first model will be used to schedule the pre-processing orders while the second model will be used to schedule the packaging orders. Constraints (6.12) will be enforced through a set of constraints that link the solution of the make model to constraints in the packaging model. This solution approach will be explained further on in this section.

7.1.2 Standard and preferred production line assignments

The production schedule that is currently used at the case-company is based on a standard fixed machine assignment and production sequence for each order. This practice often results in unfavorable schedules that may even be infeasible if they are not changed by an experienced team leader along the way.

However, an advantage of the standard routing is that the machine assignment for each order is actually the preferred assignment for that order. This preference is based on two aspects of the process.

The first of these aspects is the ease of packaging a certain order on a certain production line. Although work orders can often be produced on an array of different machines, one machine may be able to produce these products at a higher quality or at a slightly greater rate because it is specialized for that type of product.
The second aspect is based on the experience of the operator that operates a certain work center. Operators are usually assigned to work centers on which they have prior experience. This means that they know every recipe of the products that are normally assigned to that line by heart. If an operator has to produce an order that is not normally assigned to his production line, he might have to look up the recipe before he can start. This will delay the production process and result in a longer processing time for the order. Furthermore, the operator may not notice production flaws because he has little experience with how the product should look.

Both these aspects of the process result in a desire to limit the amount of orders that are not assigned to their standard production line. In order to include this desire into the model presented in section 6, the objective function would have to be extended. A penalty would be introduced that penalizes the assignment of orders to non-preferred production lines. This would make the model more complex and harder to steer because a fourth penalty would have to be introduced.

The sequential solution approach that is presented in this section will not introduce an extra penalty to the objective function. The approach will make use of the standard preferred machine of each order to find an initial solution to the scheduling problem. In this solution, all orders are assigned and fixed to their standard routing. This may result in model solutions that are infeasible because the schedule does not fit in the available production time. However, all orders will be assigned to their preferred machine.

After this initial solution is generated, the solution can be iteratively improved by relaxing the variables that were fixed earlier. This will get rid of the infeasibilities in the solution but because only a limited amount of changes will be made to achieve this, the major part of the orders will still be assigned to their preferred work centers.

The preferred machine assignment of orders will thus be used to easily generate an initial starting solution. This starting solution will, at the same time, ensure that a high number of orders will be assigned to their preferred work centers in the final solution.

7.2 A mathematical model for the pre-processing orders

To reduce the size of the problem instance, the scheduling problem will be split. This will result in a separate problem instance for each of the four departments at Heemskerk. These problem instances will only contain orders that are assigned to that department. As was explained in the previous section, in order to satisfy the constraints that connect the pre-processing orders to the packaging orders (6.12), a separate model has to be created for the pre-processing and packaging orders.

For completeness sake both models will be listed completely here. The changes that were made from the original model will be shown and explained as well.

7.2.1 Mathematical formulation

Sets

W	Batches / Work Orders
$W^m \subset W$	Make batches (batches that produce intermediates)
$W_w^m \subseteq W^m$	Make batches that supply pack batch $w \in W^p$
$W_w^r \subseteq W^m$	Make batches required to be finished before make batch $w \in W^m$
J	Work centers (production units)
$J_w \subseteq J$	Work centers that can produce batch $w \in W$
$J_{ww'} \subseteq J$	Work centers that can produce both $w \in W$ and $w' \in W$

Parameters

α_w	Processing time per unit of batch $w \in W$
δ_w	Time required before batch $w \in W$ can be used in another batch
$\omega_{ww'}$	Penalty for a switchover between batch $w \in W$ and batch $w' \in W$
β_w	Number of units required of batch $w \in W$
M	A sufficiently large number
D_w	Preferred time at which all work centers should have completed their batches

Continuous Variables

S_w	Start time of batch $w \in W$
D_w^+	Amount of time that batch $w \in W^p$ finishes after the preferred completion time D, i.e., the tardiness of order w

Binary Variables

F_{wj}	1, if batch $w \in W$ starts the processing sequence on workcenter $j \in J_w$
X_{wj}	1, if batch $w \in W$ is assigned to but not the first to be processed on workcenter $j \in J$
$Y_{ww'}$	1, if batch $w \in W$ is processed right before batch $w' \in W$

Objective function

$$\boldsymbol{min\ \ p_1 \sum_{w \in W} \sum_{w' \in W} \omega_{ww'} \cdot Y_{ww'} \ + p_2 \sum_{w \in W} S_w + p_3 \cdot \sum_{w \in W} D_w^+} \tag{7.1}$$

Allocation constraints

$$\textstyle\sum_{w \in W_m} F_{wj} \leq 1 \quad \forall \ j \in J \tag{7.2}$$
$$\textstyle\sum_{j \in J} F_{wj} + \sum_{j \in J} X_{wj} \ \ = 1 \quad \forall \ w \in W \tag{7.3}$$

Sequencing-allocation matching constraints

$$F_{wj} + X_{wj} \leq X_{w'j} - Y_{ww'} + 1 \quad \forall (w \in W^m \ ; \ w' \in W^m \ ; \ j \in J_{ww'}) \tag{7.4}$$
$$Y_{ww'} + Y_{w'w} \leq 1 \ \forall \ (w \in W^m \ ; w \in W^m) \tag{7.5}$$
$$F_{wj} + X_{wj} \leq 1 - Y_{ww'} \ \forall \ (w \in W^m \ ; w' \in W^m ; j \in (J_w \setminus J_{ww'})) \tag{7.6}$$

Sequencing constraints

$$\textstyle\sum_{j \in J_w} F_{wj} + \sum_{w\prime} Y_{w\prime w} = 1 \quad \forall \ w \in W^m \tag{7.7}$$
$$\textstyle\sum_{w' \in W} Y_{ww'} \leq 1 \quad \forall \ w \in W^{\boldsymbol{m}} \tag{7.8}$$

Timing constraints

$$S_{w'} \geq S_w + \ \beta_w \alpha_w + \delta_w \quad \forall \quad (w' \in W^m \ ; w \in W_{w'}^r) \tag{7.9}$$
$$S_{w'} \geq S_w + \beta_w \alpha_w - M(1 - Y_{ww'}) \ \forall \ (w \in W^m \ ; w' \in W^m) \tag{7.10}$$

Additional constraints

$$S_w \ \ \leq D + D_w^+ \quad \forall \ w \in W^m \tag{7.11}$$
$$F_{w,j} + X_{w,j} = 1 \ \forall \ (w \in W \ , j \in J_w) \tag{7.12}$$

Variable constraints

$$S_w \geq 0 \quad \forall \ w \in W^m \tag{7.13}$$
$$D_w^+ \geq 0 \quad \forall \ w \in W^m \tag{7.14}$$
$$X_{wj} \in \{0,1\} \ \forall \ (w \in W^m \ ; j \in J_w \ ) \tag{7.15}$$
$$Y_{ww'} \in \{0,1\} \ \forall \ (w \in W^m \ ; w' \in W^m) \tag{7.16}$$
$$F_{wj} \in \{0,1\} \ \forall \ (w \in W^m \ ; \ \ j \in J_w) \tag{7.17}$$

7.2.2 Changes made to the original model

This model will be used to schedule the orders of the pre-processing department. Therefore the set W^p is not included because it contains only pack batches.

Although the preferred finishing time D_w for each order w is still included, it is used in a different way here. Here, the preferred finishing time is no longer set by the time at which the last shift ends but it is the time at which the intermediate is required to be ready for packaging into the final product. Since each intermediate product may be required at a different time. This time is dependent on w.

This means that when product w' is required in product w and the solution of the packaging problem shows that product w starts production at time T_w then the preferred delivery time $D_{w'}$ of product w' is equal to $T_w - \delta_w$.

By setting the preferred delivery time in this way, the solution of the pre-processing model can be linked to the solution of the packaging model. Because D_w will be lower in this case than in the original case the objective function will still minimize the amount of orders that finish after the last shift ends as well. This will therefore not result in negative side effects for the scheduled finishing times of the orders.

In this model the constraints that are associated with the cross-contamination restrictions (6.10) are excluded because they are only valid for packaging orders. Furthermore constraints (6.12), which involve the connection between pre-processing and packaging orders, are also excluded because the model will contain no packaging orders.

The last constraints that are not included in this version of the model are (6.14) because they are specific to the packaging orders as well.

7.3 A Mathematical model for the packaging orders

A separate model was created to solve the scheduling problem for each of the three packaging departments at the case-company. This model differs slightly from the original model and the model for the pre-processing stage. It is listed here along with an explanation of the changes.

7.3.1 Mathematical formulation

Sets

W^m	Make batches (batches that produce intermediates)
W^p	Pack batches (batches that produce final products)
$W_w^m \subseteq W^m$	Make batches that supply pack batch $w \in W^p$
J	Work centers (production units)
$J_w \subseteq J$	Work centers that can produce batch $w \in W^p$
$J_{ww'} \subseteq J$	Work centers that can produce both $w \in W^p$ and $w' \in W^p$

Parameters

α_w	Processing time per unit of batch $w \in W^p$
δ_w	Time required before batch $w \in W^p$ can be used in another batch
$\omega_{ww'}$	Penalty for a switchover between batch $w \in W^p$ and batch $w' \in W^p$
β_w	Number of units required of batch $w \in W^p$
M	A sufficiently large number
D	Preferred time at which all work centers should have completed their batches
G_w	Starting time of Make Batch *w* according to the solution of the make batches
$K_{ww'}$	Binary matrix that is equal to 1 when for a pair of batches *w* and *w'* the variable $Y_{ww'}$ should be equal to 0.

Continuous Variables

S_w	Start time of batch $w \in W$
D_w^+	Amount of time that batch $w \in W^p$ finishes after the preferred completion time D, i.e., the tardiness of order w

Binary Variables

F_{wj}	1, if batch $w \in W$ starts the processing sequence on workcenter $j \in J_w$
X_{wj}	1, if batch $w \in W$ is assigned to but not the first to be processed on workcenter $j \in J$
$Y_{ww'}$	1, if batch $w \in W$ is processed right before batch $w' \in W$

Objective function

$$\boldsymbol{min\ p_1 \sum_{w\in W}\sum_{w'\in W} \omega_{ww'} \cdot Y_{ww'} + p_2 \sum_{w\in W} S_w + p_3 \sum_{w\in W} D_w^+} \tag{7.18}$$

Allocation constraints

$$\sum_{w\in W^p} F_{wj} \leq 1 \quad \forall\ j \in J \tag{7.19}$$

$$\sum_{j\in J} F_{wj} + \sum_{j\in J} X_{wj} \ = 1 \quad \forall\ w \in W^p \tag{7.20}$$

Sequencing-allocation matching constraints

$$F_{wj} + X_{wj} \leq X_{w'j} - Y_{ww'} + 1 \quad \forall (w \in W^p\ ;\ w' \in W^p\ ;\ j \in J_{ww'}) \tag{7.21}$$

$$Y_{ww'} + Y_{w'w} \leq 1\ \forall\ (w \in W\ ; w \in W) \tag{7.22}$$

$$F_{wj} + X_{wj} \leq 1 - Y_{ww'}\ \forall\ (w \in W\ ; w' \in W; j \in (J_w \setminus J_{ww'})) \tag{7.23}$$

Sequencing constraints

$$\sum_{j\in J_w} F_{wj} + \sum_{w'} Y_{w'w} = 1 \quad \forall\ w \in W^p \tag{7.24}$$

$$\sum_{w'\in W} Y_{ww'} \leq 1 \quad \forall\ w \in W^{\boldsymbol{p}} \tag{7.25}$$

$$Y_{ww'} = 0 \quad \forall (w \in W^p\ ; w' \in W^p)\ |\ K_{ww'} = 1 \tag{7.26}$$

Timing constraints

$$S_{w'} \geq G_w + \delta_{w'} \quad \forall\ \ (w' \in W^p\ ; w \in\ W_{w'}^m) \tag{7.27}$$

$$S_{w'} \geq S_w + \beta_w \alpha_w - M(1 - Y_{ww'})\ \forall\ (w \in W^p\ ; w' \in W) \tag{7.28}$$

$$S_w \geq 1 \quad \forall\ (w \in W^p) \tag{7.29}$$

Additional constraints

$$S_w + \alpha_w \beta_w \leq D + D_w^+ \quad \forall\ w \in W^p \tag{7.30}$$

$$F_{wj} + X_{wj} = 0\ \forall\ (w \in W^p\ ; J \notin J_w) \tag{7.31}$$

$$Y_{ww'} = 0\ \forall\ (w \in W^p; w' \in W^p; j \notin J_{ww'}) \tag{7.32}$$

Variable constraints

$$S_w \geq 0 \quad \forall\ w \in W \tag{7.33}$$

$$D_w^+ \geq 0 \quad \forall\ w \in W \tag{7.34}$$

$$X_{wj} \in \{0,1\}\ \ \forall\ (w \in W\ ; j \in J_w\ ) \tag{7.35}$$

$$Y_{ww'} \in \{0,1\}\ \forall\ (w \in W\ ; w' \in W) \tag{7.36}$$

$$F_{wj} \in \{0,1\}\ \forall\ (w \in W\ ;\ \ j \in J_w) \tag{7.37}$$

7.3.2 Changes made to the original model

The separate model that will be used for the packaging orders of each of the three packaging departments slightly differs from the original model presented in section 6.

Although this model is only used for the scheduling of pack batches, the sets of make batches W^m and W_w^m are included as well. These will be used to enforce the constraints that link the starting times of make batches to the starting time of the pack batches they are required in. The subset of make batches W_w^r is no longer included because it is only required to schedule the make batches.

The set of parameters has been extended by the parameter G_w. This parameter can be used to store the starting times of all pre-processing batches in the set W^m. This parameter will be used to connect this model to the model for make batches (section 7.2) after it has been solved. The parameter will be provided when the solution from a certain instance of the make model is connected to the model presented here.

To connect the pre-processing batches to the pack batches, constraints (6.12) are altered. Because the model for the make batches is solved separately, the starting times S_w of these batches is not available. The starting time will be provided in the form of parameter G_w. The resulting constraints (7.27) ensure that a packaging batch w' starts at least $\delta_{w'}$ after the starting time of its required make batches $w \in W_{w'}^m$.

Through these constraints a solution generated using the model for pre-processing orders can be connected to this model for the packaging orders. This ensures that the complete solution is still feasible. This will be used in the sequential approach that is presented in the next section.

From here on, the approach presented in this section, i.e. sequentially solving different parts of the model in 3 distinct steps will be referred to as the ‘sequential approach’.

7.4 A sequential solution approach to solve the scheduling problem

Until now three different optimization models have been presented. The first presented model in section 6 is an optimization model that can be used to schedule an entire instance of the scheduling problem at once. Unfortunately, as was shown in section 6.4, using this method, the problem size of one production day is too large for a commercial solver to efficiently handle.

To be able to find a good solution in a reasonable amount of time, the problem instance was split up. After this separation each of the four instances of the separate processing departments will be solved independently of each other. In order to be able to create feasible schedules using this approach, two different versions of the original model were presented. The model presented in section 7.2 will be used to solve the model for the pre-processing department while the model presented in section 7.3 will be used for all three instances of the packaging department.

Here, a step-by-step explanation of the used approach will be given. Because the connection between pre-processing orders and packaging orders needs to stay feasible, the models for these instances will be solved sequentially. In each step the solution of the previous step will be used as input for the model in the next step. It should be noted that when it is indicated that a model instance is 'solved', the commercial solver *CPLEX 12.5* is used to find a solution for the presented problem instance.

The approach will start at the back of the production process. All packaging orders will be scheduled first. After this, the solution to this problem will be used as input to solve the model for the pre-processing orders. This choice was made because the packaging process is considered to be the most important process. During the packaging of products, most of the value is added to the product. The production of intermediates is also faster than the time it takes to package the intermediates. This makes the packaging process the bottleneck. Lastly, idle time on packaging work centers is generally considered a bad thing because personnel are not able to go do something else. The operator of a packaging line is often only trained to operate that specific work center. In the pre-processing department, idle time is of less concern. When workers there have to wait before they can continue processing on their line, they can do other things like cleaning, processing pre-work for tomorrow, or helping out their colleagues on other lines.

The sequential approach employs the following steps to solve the model for all orders of a single production day:

1. In the first step, the three different instances for the three separate packaging departments are solved independently using the model presented in section 7.3. In this step the instances are solved without the connection to their respective required pre-processing orders ($G_w = 0 \ \forall \ w$).

   When G_w equal 0 for all orders, constraints (7.27) are effectively removed from the problem. This holds because $S_{w'} \geq 1$ and $\delta_{w'} = 0{,}5$ for all packaging orders. This means that packaging orders can be scheduled anywhere in time without taking into account the time at which their required intermediates become available. This passes with the fact that a good schedule for the packaging orders is considered most important and that these orders should at first be allowed to be scheduled anywhere in time.

   The three different problem instances for the bag, bowl and high care packaging departments can be solved independently and in parallel because none of the orders in these departments can be assigned to work centers in the other departments.
   To further increase the speed at which the three model instances can be solved. The following steps will be performed during the solving of the problem instances:

a. As explained in section 7.1.2, all orders have a preferred work center assignment. In the first iteration, the model will be solved while all orders are fixed to their preferred work centers. This means that F_{wj} and X_{wj} are fixed to 0 for all combinations of w and j that are not preferred order-work center combinations.

   Furthermore, all changeover variables $Y_{ww'}$ of combinations of orders w and w' that do not have the same preferred machine assignment, are fixed to 0 as well. After all it is not possible for a changeover to occur from order w to order w' when both are assigned to different work centers.

   Because any fixed variables can be treated as parameters, this greatly reduces the number of variables in the model and allows for much quicker solving during this step.

b. After step 1a has been performed, the solution may not be very good. Because orders are fixed to their preferred machines, the entire workload is assigned to that production line. If the total workload is larger than what is available during the two available shifts, an unfavorable and unfeasible solution will be the result. In this case the third part of the objective function (associated with p_3) will be greater than 0.

   In order to arrive at a feasible and/or favorable solution, some of the workload on the overburdened production lines will have to be assigned to alternative production lines that still have capacity available. This will be done in this step.

   The variables that were fixed in step 1a are now selectively unfixed. The variables F_{wj} , X_{wj} and $Y_{ww'}$ that denote orders that are assigned to overburdened machines are unfixed. This means that these orders can be freely assigned to any work center they can be processed on.

   After this, the model is solved again until an acceptable solution is found in which production lines are no longer overburdened. Because only orders that are assigned to overburdened production lines are allowed to switch over to another production line, only a limited amount of orders will be assigned to non-preferred work centers. This passes with the objective of having as many preferred work center assignments as possible.

An illustrative but over exaggerated example of what happens during step 1a and 1b is given in figure 7.1 and 7.2 respectively.

In this example, the preferred production line for all orders is the HC01. Naturally, all orders are thus assigned to this production line in step 1a. Because the shift of this production line normally ends at $D_w = 20$, this schedule finishes long after the last shift ends and is therefore infeasible. In step 1b the variables of the orders on line HC01 are relaxed (in this case all orders). The result is shown figure 7.2. It can be seen that after this step, the complete schedule finishes within the set timeframe. The penalty in the objective function that is associated with p_3 is reduced to 0 and the overburdened workload has been moved to alternative packaging line HC02.

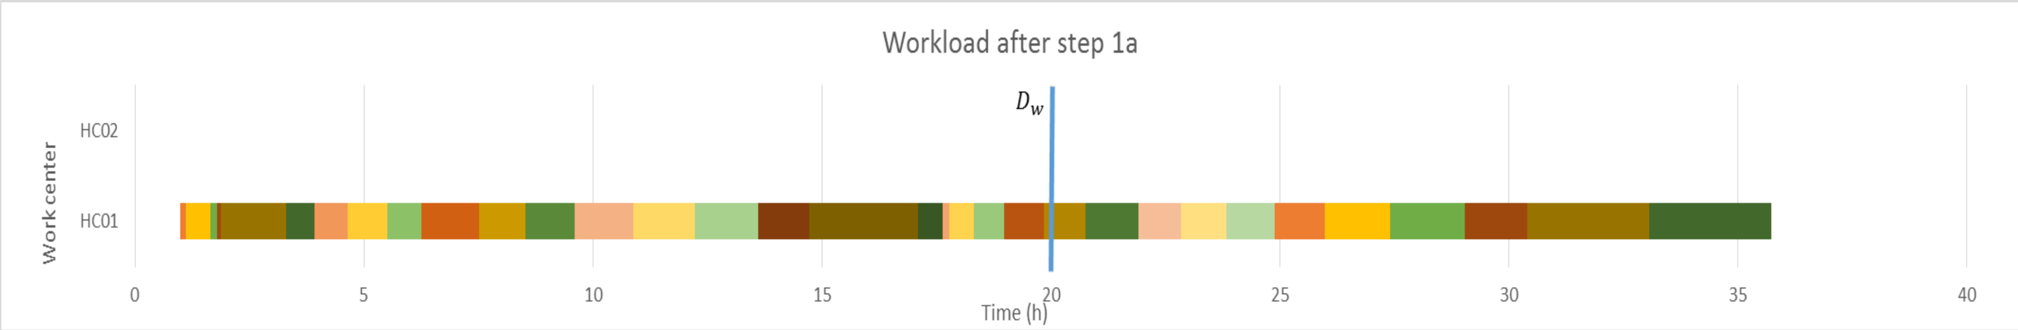

Figure 7.1: Schedule for order that have preferred work center HC01. Result after step 1a of the sequential approach.

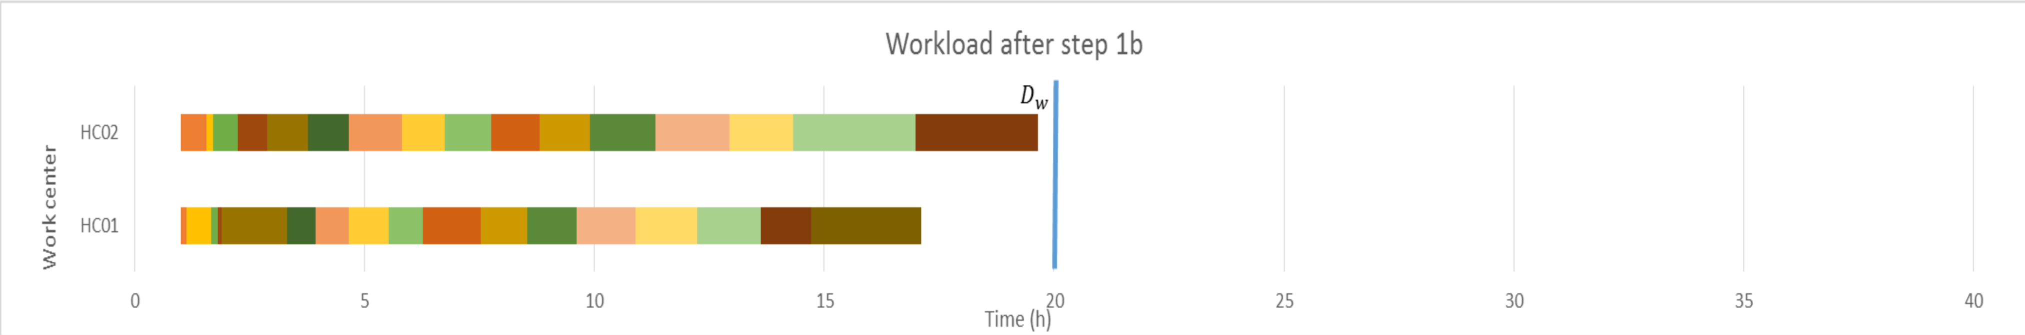

Figure 7.2: Schedule for order that have preferred work center HC01. Result after step 1b of the sequential approach.

2. In step 1 the initial solution for the three instances of the packaging departments was found. In step 2 the solution for the pre-processing department will be generated. Because the final solution will have to be feasible, the constraints that link the pre-processing department to the packaging departments are already included to some extent.

   From the solution for the packaging departments, the earliest times at which each intermediate is required can be calculated. If a certain intermediate w is required in three different packaging orders, then the packaging order $w'$with the earliest starting time is leading. If the intermediate available on time for this order, then it will be on time for the other two orders as well.

   To connect the pre-processing order to its related packaging orders, the original model uses constraints (6.12)

   These constraints ensure that a packaging order w' starts at least $\delta_{w'}$ after production of the intermediates it requires have started. Since the starting times of the packaging orders are known from the solutions in step 1, the preferred starting times of the intermediates are known as well. These times are equal to $S_{w'} - \delta_{w'}$ for all of the intermediates required in packaging order w'.

   In this step, the relationship between pre-processing and packaging orders is not enforced through a constraint but through an incentive in the objective function. This is done by setting the preferred delivery time D_w of each of the orders to their preferred starting time (as calculated above). The amount of orders that starts after this preferred starting time is then minimized in the objective function. If an order finishes after its preferred starting time D_w it is given a penalty in the objective function.

   In this way, the number of orders that result in an infeasible intermediate-final product connection is minimized. This number can be steered through penalty factor p_3. A higher penalty factor will result in a higher number of orders that satisfy constraints (6.12).

   Summarizing step 2 of the solution approach:

   a. Preferred starting times for all intermediates are calculated from the solutions found for the packaging problems in step 1. This is done by setting:

      $$D_w = S_{w'} - \delta_{w'} \quad \forall \ (w' \in W^p \, ; w \in W^m_{w'})$$

   b. The calculated D_w are then used to solve the model for the pre-processing orders. This model was presented in section 7.2.

   c. Penalty factor p_3 can be used to steer the amount of orders that are scheduled 'on-time' for usage in their respective final products.

It should be noted that in solving the model for pre-processing batch; it is unnecessary to fix or unfix variables (like in step 1a and 1b). Almost all pre-processing orders have only one allowed machine assignment. Therefore almost all assignment variables have only one option for assignment.

3. In step 1 and 2 of the solution approach, solutions have been found for all orders in the pre-processing and packaging departments. Although these solutions are feasible in their respective departments, the total solution may not be feasible.

   The total solution may be infeasible because not all pre-processing orders are scheduled to start at an appropriate time during step 2 of the approach. This will result in packaging orders that start before their required intermediates are ready for consumption.

   To arrive at a uniformly feasible solution, one of two options can be chosen:

   1. Use the available capacity at the pre-processing department to schedule the pre-processing batches that result in infeasibilities on the day before they are required. This is known as pre-work.
   2. Adjust the schedules of the packaging departments to the solution for the make batches found in step 2.

   When only a small amount of pre-processing batches results in infeasibilities, choosing option 1 may be an easy way to arrive at a feasible schedule. If all these batches can be scheduled on the day before, the schedule for the packaging orders does not have to be adjusted.

   In the current situation at Heemskerk it is common practice in the pre-processing department to perform some 'pre-work' for the next day in the off-hours or when the orders of the current day are finished. There is often plenty of time available for this pre-work. Therefore it would not be hard to adapt option 1 and select a number of the orders as obligatory pre-work.

   When the number of infeasibilities in the overall solution is high or when not all infeasible orders can be executed as pre-work, choosing option 1 may not be viable. In this case, the solution of the packaging departments will have to be adjusted. This is option 2.

   For option 2, constraints (7.27) are used to introduce the connection between the pre-processing and packaging orders into the model for the packaging departments.

   In this step, G_w is set equal to the starting times of the make orders that were found when solving the model in step 2. Constraints (7.27) then ensure that packaging orders can only start production after their intermediates are available. Solving this model for the packaging departments thus ensures a feasible solution for the entire problem.

   Subsequently the model for the packaging departments is solved again with constraints (7.27) included. For this, the same approach is used as in step 1a and 1b.

   It is also possible to use a combination of option 1 and 2 where a part of the infeasible orders are scheduled as pre-work and the remaining infeasibilities are resolved by adjusting the packaging solution. This is essentially the same as choosing option 2 because the instances of the packaging departments still have to be resolved.

Summarizing step 3 of the solution approach:

a. Step two of the approach may result in an infeasible overall schedule. One of two options can be chosen to arrive at a feasible solution:
   i. Schedule the infeasible pre-processing orders as pre-work on the day before. This effectively removes the orders and their infeasibilities from the problem.
   ii. Adjust the schedule of the packaging department to the solution found in step 2 (or the remaining infeasibilities after partially rescheduling some orders as pre-work).

b. If option ii is chosen, the problem instances of the three packaging departments have to be solved again. Here, constraints (7.27), based on the solution found in step 2, ensure an overall feasible solution.

Ultimately it is up to the management of the company to decide which approach should be taken in step 3 of the approach. Option 1 may result in more changeovers on the pre-processing department. On the other hand option 2 may result in a worse solution for the packaging department because they will have to wait for their intermediates to be ready.

For a good comparison, both approaches for step 3 that were listed here will be discussed in the results section. These can be found in section 8.

Because the solution approach is rather complex and involves a lot of different aspects, a visual representation of the process is given in figure 7.3. The denotations used are equal to the ones used in the textual description (Step 1, 2, etc.).

7.4.1 Iterations

The approach presented in this section solves the problem sequentially. Different parts of the problem are solved in sequence and independent of each other. Due to this independence, it may be possible that the overall solution can be improved by iterating between the different parts of the problem. This would mean solving the instance in step 2, then solving step 3 and then go back to the instance of pre-processing batches from step 2 to see if this can somehow be adjusted to give a better solution for the packaging orders.

There are several challenges that make it hard to introduce these iterations into the current approach. First, it would require introducing an extra incentive into step 2 of the approach. The few pre-processing orders that still result in infeasibilities after step 2 would have to be given an extra incentive, e.g. extra individual weights for the third part of the objective function, to be scheduled on time. This may then result in other orders that are scheduled late. These, in turn, will require further adjustment of schedule for the packaging orders.

A second challenge is that the variable S_w used to minimize the idle time is not an ideal candidate for this job. In section 8.2.1 it will be shown that a large part of the sum of this variable cannot be removed without introducing overlap, and thus infeasibility, in the schedule. When, for example, the total sum of all starting times is 800 hours and the total idle time is 3 hours. Then the second part of the objective function could only be reduced by a small amount before the ideal situation is achieved (no idle time). This would only have a small effect on the total value of the objective function. The result of this, as will be shown in section 8, is that it is hard to limit the idle time during step 3. When iterations are introduced to the approach, step 3 would have to be performed multiple times. In each iteration the amount of idle time would increase. To prevent this, a different way of limiting the idle time would have to be found. This will however require extensive adjustments of the model.

Furthermore, the schedule for the packaging orders that is found during step 1b can be viewed as the most ideal schedule for these orders that can be found using this approach. By solving the problem for the packaging orders first and then solving the problem for the pre-processing orders in step 2, the packaging orders have more freedom in the scheduling process. This is exactly what this approach is trying to achieve. As a result of this unbalance, each iteration between step 2 and 3 will worsen the schedule of the packaging orders found in step 1b. Introducing iterations would thus go against one of the objectives of the sequential approach.

Lastly, it is unlikely that introducing iterations will show a great improvement in the overall solution. As will be shown in a comparison between the MIP-based approach and the sequential approach in section 8.2, the difference between the solutions of these two approaches is rather small.

Because of the above, iterations were not performed between step 2 and step 3.

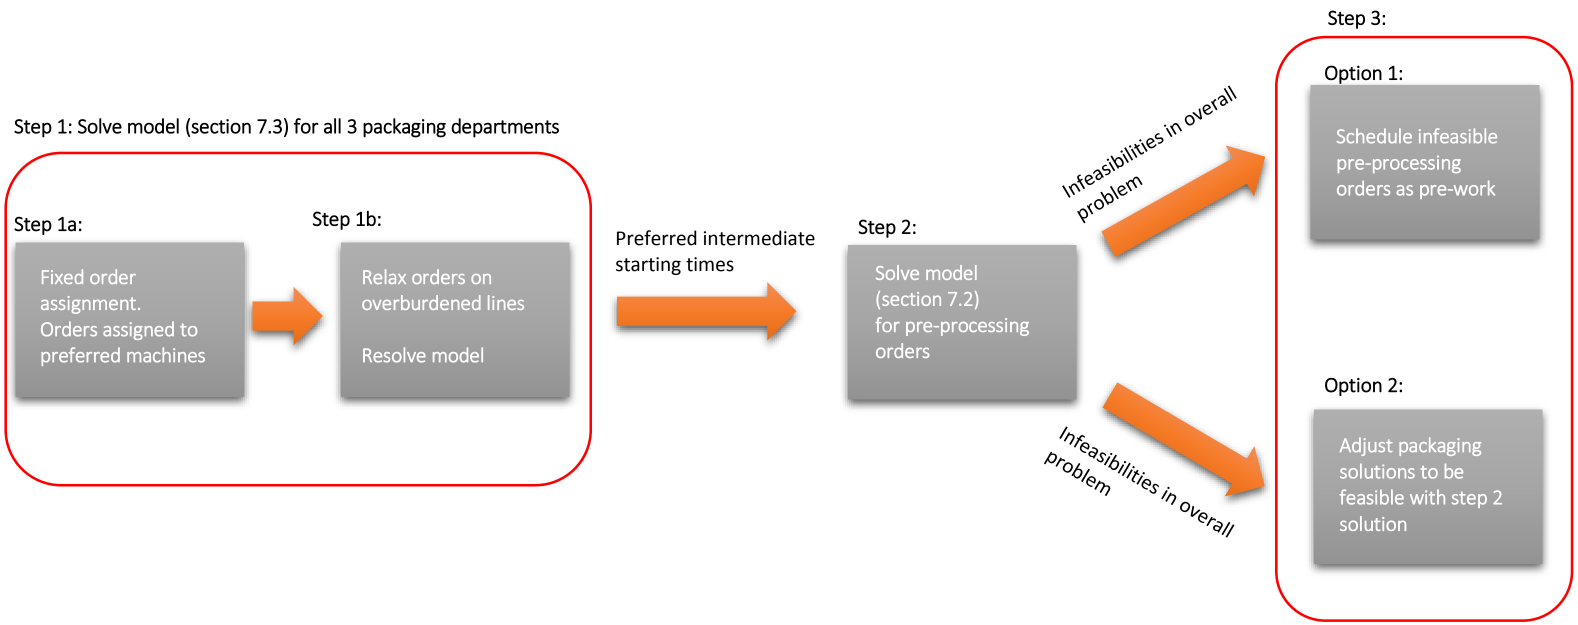

Figure 7.3: Visual representation of the sequential approach used to solve the production-scheduling problem at Heemskerk

8 Model results and comparison

So far, two different approaches were presented to solve the problem. The first approach involves solving a single optimization model for the complete problem instance of a single production day. It was presented in section 6. The preliminary results for this approach showed that finding a feasible solution in a reasonable amount of time using a commercial solver was unlikely.. Therefore, in section 7, a sequential approach was presented to find a solution for the problem. In this section, a number of exemplary results will be presented to illustrate the way in which a schedule is generated by the models. Furthermore, the approach from section 6 will be compared to the sequential approach from section 7. The solution will also be compared to the current situation at Heemskerk. This should give some insight in the possibilities of the approach. All solutions were found using an Intel Core i5-3317U CPU and 8GB DDR3 ram.

8.1 A toy problem to illustrate the process.

In order to verify the correct working of the model and to illustrate the results, a toy problem was devised. This problem includes a few exemplary orders to show the different aspects of the process.

In the toy problem, orders for the packaging of two different types of cut vegetables and their intermediates are scheduled. Different packaging sizes and intermediates result in a total of 10 packaging orders and 4 intermediate orders that need to be scheduled. Table 8.1 lists all the orders and their properties.

Table 8.1: Orders and details for a toy problem of the production process.

Order nr.	*Description*	*Amount required*	*Time per unit (m)*	*Available work centers*	*Required intermediates*	*Dirt Level*	*Type of product*
1	Andijvie 600gr	1688 ce	0,0444	B01, B04	13	3	Final
2	Soepgroenten 500gr	1550 ce	0,0208	B03		3	Final
3	Andijvie 400gr	3455 ce	0,0336	B01, B04	13	2	Final
4	Soepgroenten 150gr	1342 ce	0,0180	B03	14	1	Final
5	Sal. Andijvie 400gr	7250 ce	0,0299	B01, B04	13	2	Final
6	Andijvie fijn 200gr	3145 ce	0,0343	B01, B04	11	1	Final
7	Sal. Soepgroenten 300gr	15300 ce	0,0140	B03	14	2	Final
8	Soepgroenten fijn 200gr	1008 ce	0,0198	B03	14	1	Final
9	Andijvie fijn 450gr	1450 ce	0,0547	B01, B04	11	1	Final
10	Soepgroenten 300gr	3008 ce	0,0164	B03	14	2	Final
11	Andijvie sliert 4mm	1333 kg	0,1098	A09	-	1	Intermediate
12	Winterpeen blok 8mm	726 kg	0,0281	C01	-	1	Intermediate
13	Andijvie sliert 8mm	5698 kg	0,0501	A09	-	1	Intermediate
14	Soepgroenten Mengsel	6670 kg	0,0191	A01	12	1	Intermediate

This toy problem is forged in such a way that a lot of aspects of the solution are known in advance. All curly endive ("Andijvie" in the table) should be scheduled to either the B01 or B04 production line while all soup vegetables ("Soepgroenten" in the table) should be assigned to machine B03. Furthermore, the dirt levels are set up in such a way that the largest packaging weights should be scheduled last. Lastly the preferred delivery time D_w for the packaging orders was set equal to 5 hours. Because the total endive workload is 9,92 hours, the model should distribute this work between production line B01 and B04.

The model presented in section 6 was translated using the mathematical modeling software suite *AIMMS*. Subsequently, the commercial solver *CPLEX 12.5* was used to solve the model using the main data presented in table 8.1.

The solution that was found is presented below. Table 8.2 gives a summary of the results found with the model while table 8.3, 8.4, 8.5 and 8.6 respectively give the variables F_{wj}, X_{wj}, $Y_{ww'}$ and S_w. Figure 8.1 shows a visual representation of the schedule.

Table 8.2: Summary of the solution found after solving the model for the toy problem

Description	*Value*
Used penalty factor (p_1, p_2, p_3)	(1,.01,10)
Best Solution (Objective value)	82.92
Best LP bound	82.92
% Gap	Optimal solution
Solution time (s)	0.27 s
# of Constraints	719
# of Variables	326
# of orders that are scheduled late (Finishing time > D_w)	3
Total amount of idle time (h)	1,18 (h)
Total changeover penalty	28

Because the solved problem is a toy problem, the figures presented in table 8.2 do not say much. What can be concluded is that the model solves problem instances of this small size very quickly and that it is possible to find optimal solutions. The results also show that the model correctly minimized the different penalties in the objective function.

Table 8.3: Values of variable $\boldsymbol{F_{wj}}$ after solving the model for the toy problem

w	j	value
Andijvie fijn 200gr	B01	1
Soepgroenten fijn 200gr	B03	1
Andijvie fijn 450gr	B04	1
Andijvie sliert 4mm	A09	1
Winterpeen blok 8mm	C01	1
Soepgroenten Mengsel	A01	1

The variable F_{wj} in table 8.3 is equal to 1 when an order w is the first to be processed on production line *j*. It can be seen that the orders with the lowest dirt levels are processed first indeed. This is in concurrence with the cross-contamination restrictions that are enforced by constraints (6.10) of the model.

Table 8.4: Values of variable $\boldsymbol{X_{wj}}$ after solving the toy problem

w	j	value
Andijvie 600gr	B04	1
Soepgroenten 500gr	B03	1
Andijvie 400gr	B04	1
Soepgroenten 150gr	B03	1
Sal. Andijvie 400gr	B01	1
Sal. Soepgroenten 300gr	B03	1
Soepgroenten 300gr	B03	1
Andijvie sliert 8mm	A09	1

The variable X_{wj} is equal to 1 if order w is assigned to work center j but is not the first to be processed. As expected, the model assigns all endive to the B01 and B04 production lines and all soup vegetables to the B03 production line. This is exactly the expected result.

Table 8.5: Values of the variable $\boldsymbol{Y_{ww'}}$ after solving the toy problem

w	w'	value
Andijvie 400gr	Andijvie 600gr	1
Soepgroenten 150gr	Soepgroenten 300gr	1
Andijvie fijn 200gr	Sal. Andijvie 400gr	1
Sal. Soepgroenten 300gr	Soepgroenten 500gr	1
Soepgroenten fijn 200gr	Soepgroenten 150gr	1
Andijvie fijn 450gr	Andijvie 400gr	1
Soepgroenten 300gr	Sal. Soepgroenten 300gr	1
Andijvie sliert 4mm	Andijvie sliert 8mm	1

The variable $Y_{ww'}$ is equal to 1 when order w' is processed right after order w on the same production line. Again, it is clear from table 8.5 that the cross-contamination requirements are met. For example; *Andijvie 600gr* (dirt level 3) is processed after *Andijvie 400gr* (dirt level 2). The same can be shown for the other pairs of orders for which $Y_{ww'} = 1$.

Table 8.6: Values of starting time variable S_w after solving the toy problem

Product	*Starting time (h)*
Andijvie 600gr	4,871
Soepgroenten 500gr	6,474
Andijvie 400gr	2,939
Soepgroenten 150gr	1,672
Sal. Andijvie 400gr	2,939
Andijvie fijn 200gr	1,000
Sal. Soepgroenten 300gr	2,895
Soepgroenten fijn 200gr	1,340
Andijvie fijn 450gr	1,000
Soepgroenten 300gr	2,073
Andijvie sliert 4mm	0,000
Winterpeen blok 8mm	0,000
Andijvie sliert 8mm	2,439
Soepgroenten Mengsel	0,840

Variable S_w denotes the scheduled starting time of order w. The results in table 8.6 show that constraints (6.14) are kept. These constraints ensure that all packaging orders start at least 1 hour after the start of the first pre-processing shift.

In order to illustrate what a complete schedule looks like, figure 8.1 has been created. It shows the schedule that is found after the toy problem has been solved. The different colored bars represent the work orders that are assigned to each of the work centers. The numbers on the bars are equal to the order numbers shown in table 8.1. Furthermore, the arrows in the figure indicate that a certain intermediate is required in another intermediate or final product.

The figure shows that the resulting solution complies with how the toy problem was designed. All endive-packaging orders (1, 3, 5, 6 and 9) are assigned to either machine B01 or B04. Furthermore the workload of these orders is evenly distributed between these two machines. Lastly, the largest packaging sizes (1 and 2 for example) are scheduled last.

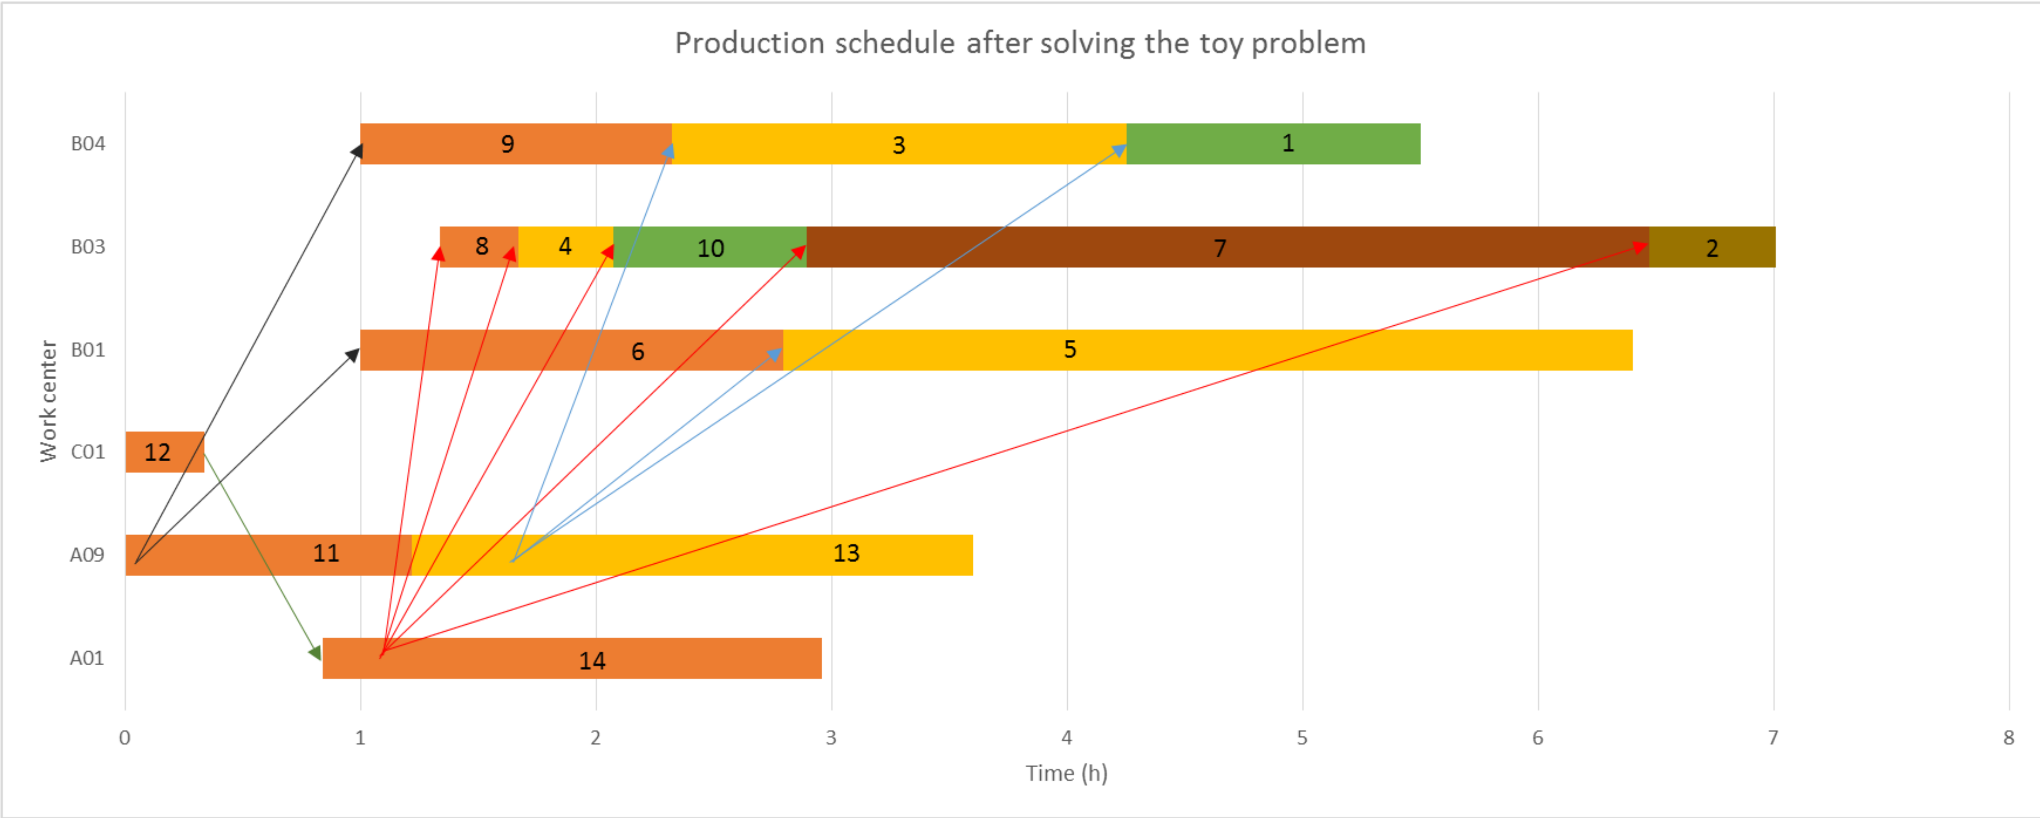

Figure 8.1: Production schedule after solving the toy problem.

8.2 Comparison of the MIP-based approach to the sequential approach.

In section 6 the complete mathematical formulation for the problem was presented. In order to be able to find a feasible solution in a reasonable amount of time, a sequential approach was presented for solving the problem in section 7. For this approach, some concessions were made to arrive at a usable approach. In this section a comparison between the two approaches is made in order to see the effects of these concessions.

To be able to solve the instance using both approaches, the size of the exemplary instance of the production day of the 11^{th} of July 2013 was reduced. This was necessary to be able to solve the problem using the MIP-based approach. Only selecting the orders that are normally assigned to 3 of the 10 bag-packaging lines reduced the problem size to a manageable level. The required intermediate orders for these packaging orders are included as pre-processing orders in this instance.

Only selecting the packaging orders from 3 of the 10 bag-packaging lines does not result in a low number of pre-processing lines that need to be scheduled. Because the required intermediates for the selected final product orders are so omnifarious, 9 of the 11 pre-processing lines would have to be included in the model. If this would not be changed, the model would be rather easy to solve. Because a lot of other intermediates would not be included, there would be very little workload on each of the pre-processing lines and scheduling each intermediate on time would be easy.

To prevent this, all pre-processing orders are converged onto two production lines. This means that in the problem instance used here, any intermediate order can be scheduled on either of two pre-processing machines. Because the total workload of all pre-processing orders is equal to about 33 hours and the available production time is 38 hours, this will result in a problem that is harder to solve. This will make the problem instance more similar to a problem instance for a full production day.

The main characteristics of the problem instance that was simulated for this test are given in table 8.6.

Table 8.6: Main characteristics of the simulated problem to compare the two solution methods

Description	*Value*
Total # of orders	60
# of packaging orders	38
# of pre-processing orders	22
Packaging lines	3 (B01, B04, B06)
Pre-processing lines	2 (A01, A02)
Total packaging workload	46.16 (h)
Available packaging time	57 (h)
Total pre-processing workload	33.09 (h)
Available pre-processing time	39 (h)

8.2.1 Results after solving the complete model

The problem instance created for this comparison is solved using the model presented in section 6 and the commercial solver *CPLEX 12.5.* A summary of the results is given in table 8.7.

It should be noted that the solver is not able to find optimal solutions for problems of this size. Therefore a stopping criterion has to be used to stop the solver. In the cases presented here, the solver was stopped after a period of 600s (10m) during which no improvement of the best feasible solution was found. All solution times presented here include these 600 seconds.

Table 8.7: Main results after solving the comparison problem using the complete model.

Description	*Value*
Used penalties in objective (p_1, p_2 *and* p_3)	(1,1,1)
Best Solution (Objective value)	**583,67**
Best LP bound	**281,72**
% Gap	51,73%
Solution time until feasible (s)	10.28 s
Total solution time	2661 s
# of Constraints	14,804
# of Variables	3,196
# of orders that are scheduled late (Finishing time > D_w)	0
Total amount of idle time (h)	2.29 h
Total changeover penalty	147
Changeover penalty for packaging orders	102
Changeover penalty for pre-processing orders	45
# of orders assigned to non-preferred packaging lines	5

The work load in both the packaging department and the pre-processing department is quite high for this problem instance. The utilizations of the departments are 80.9% and 84.8% for respectively the packaging and the pre-processing departments. Even with utilization this high, the model finds a very good solution.

The optimality gap (51.73%) is rather large. This does however not per se mean that the solution that was found is bad. The LP relaxation of the problem is just very weak. This can be shown by a simple example. Below the objective function (6.2) is shown.

$$\boldsymbol{min\ p_1 \sum_{w \in W} \sum_{w' \in W} \omega_{ww'} \cdot Y_{ww'} + p_2 \sum_{w \in W} S_w + p_3 \sum_{w \in W} D_w^+} \tag{6.2}$$

The first part of the objective function is the changeover penalty (associated with p_1). In the solution that was found this part is equal to 147. This value is the sum of 50 changeovers that each have a penalty which ranges from 1 to 10. In the highly unlikely situation that all of these changeover penalties are equal to 1, the value of the first part of the objective function would still be equal to 50.

The second part of the objective function is equal to the sum of all starting times of the orders. This part is equal to 357.43 in the solution. This number could be reduced only by reducing the amount of idle time, i.e. the time during which a production line is not used, for every production line. From table 8.7 it can be seen that the current solution already has a very small idle time. Even when the idle time could be reduced to 0, part 2 would still be equal to 330.17. Reducing this part further cannot be done without having to overlap certain orders.

Because there are no orders that finish late, part 3 of (6.2) is already equal to 0. When these numbers are put together, a 'lower bound' can be calculated for the objective function. This lower bound is equal to 380,17. This is still significantly higher than the best-found LP bound. Even when a feasible solution with this objective value would be found, the optimality gap would still be equal to 25.9%. It is highly unlikely that this solution even exists. The LP-relaxation thus gives a very weak bound.

Overall, the solution that was found is very good. The idle time on all production lines is very low at only 2.29 hours over all lines. There are no orders that finish late and there are very few changeovers that have a high changeover penalty (see table 8.8)

Table 8.8: # of occurrences of each changeover in the solution of the complete model

Individual changeover penalty	*# of occurrences of this changeover*
0	8
1	5
2	28
4	4
5	6
10	4
Total changeover penalty:	**147**

Remarkably there are also very few orders that are not assigned to their preferred work center (5 out of 38 packaging orders), even though there is no incentive in the objective function for these orders to be assigned to their preferred production line.

This phenomenon is likely the result of the minimization of bad changeovers. Because products that are normally assigned to the same work center have similar recipes and because changeovers largely depend on these recipes, the result is that by minimizing the changeover penalty, the amount of orders assigned to non-preferred production lines is also minimized.

Figure 8.2 shows the solution schedule that was found using this approach. Each differently colored block again indicates a period during which one work order is processed on that production line. The white spaces indicate that during that period, the machine is idle and no work is assigned to that line during that period.

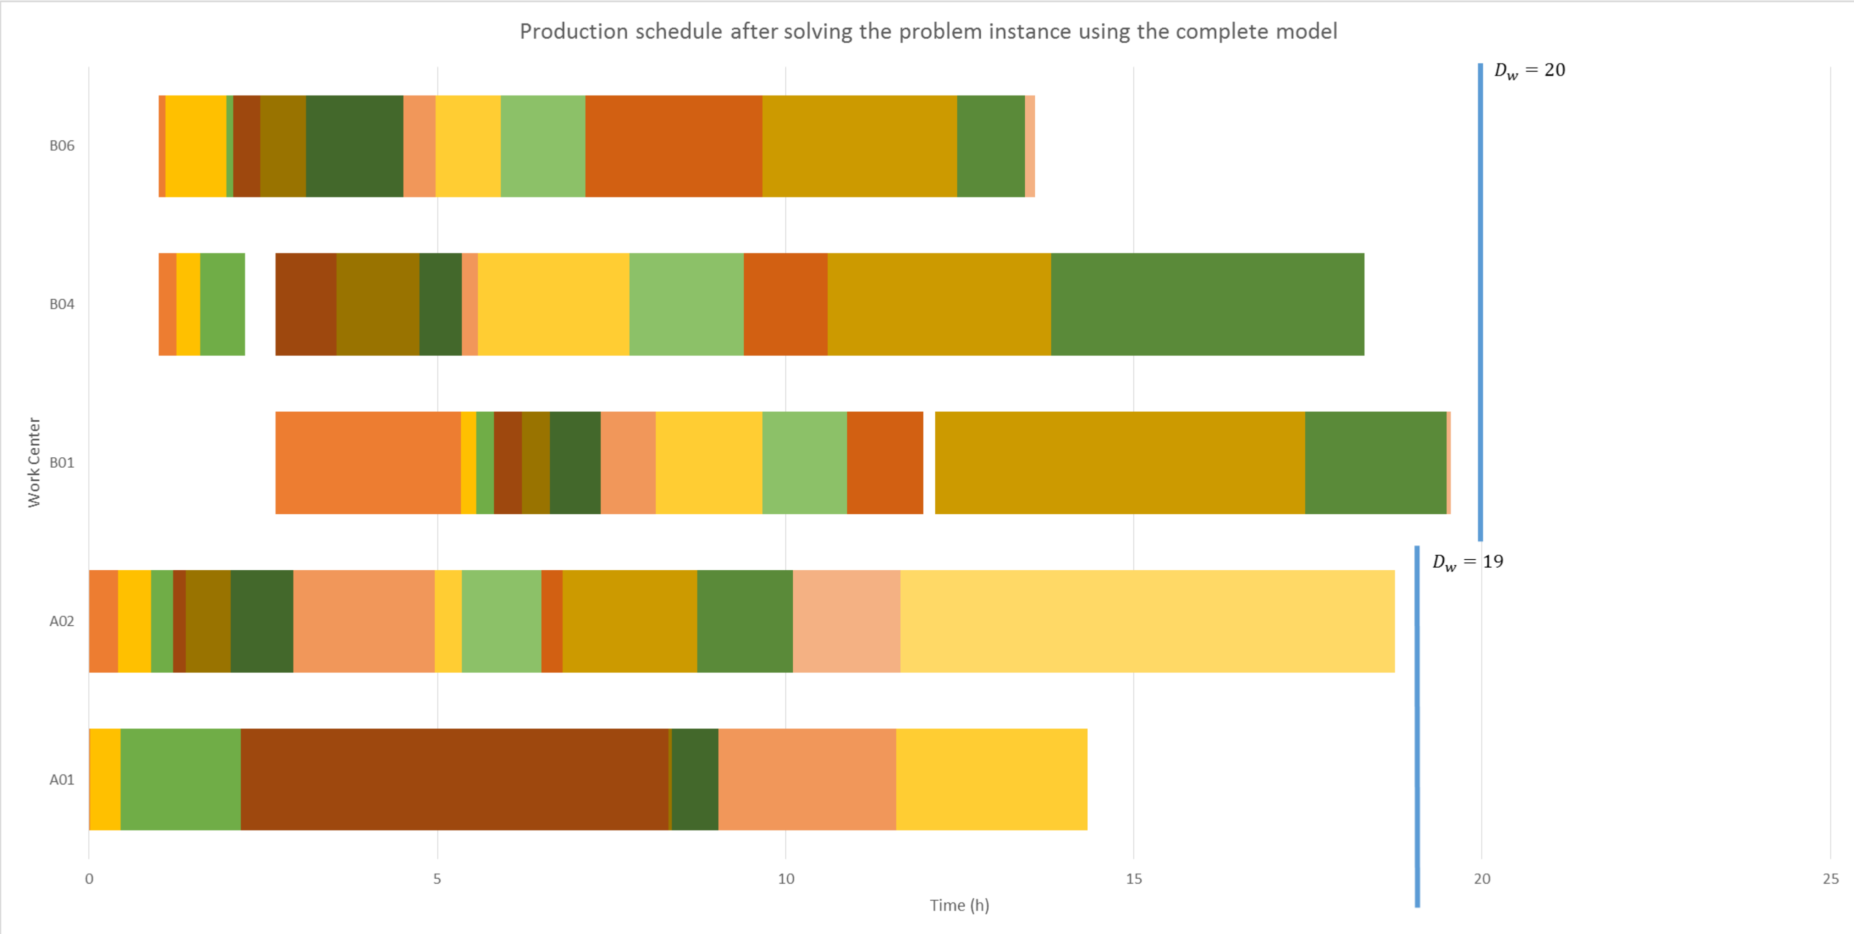

Figure 8.2: Complete production schedule after solving the problem instance using the complete model (section 6)

8.2.2 Results after solving the model using the sequential approach

For comparison, the same problem instance will now be solved using the sequential approach that was presented in section 7. Because this approach consists of several steps, each of the intermediate results will be presented and finally the complete solution will be given and it will be compared to the solution found in section 8.2.1.

Step 1a

As explained in section 7, during step 1a, the model presented in section 7.3 will be solved for all packaging orders. Because the problem instance used in this section only involves bag packaging only a single instance, instead of the three in a full instance, has to be solved. All packaging orders scheduled here, belong to the same department.

The main characteristics of the solution after step 1a are given in table 8.9. It should be noted that the results only include packaging orders in this step. They cannot be compared to the results in table 8.7 yet.

Table 8.9: Main solution characteristics of the packaging orders after step 1a of the solution approach

Description	*Value*
Used penalties in objective (p_1, p_2 and p_3)	(1,1,1)
Best Solution (Objective value)	**368,48**
Best LP bound	**216,31**
% Gap	41.46%
Solution time until feasible (s)	During pre-solve (<0.44s)
Total solution time	714s
# of Constraints	4575
# of Variables	655
# of orders that are scheduled late (Finishing time > D_w)	2
Sum of workload scheduled after shifts have finished	3.69 h
Total amount of idle time (h)	0
Total changeover penalty	96
# of orders assigned to non-preferred packaging lines	0

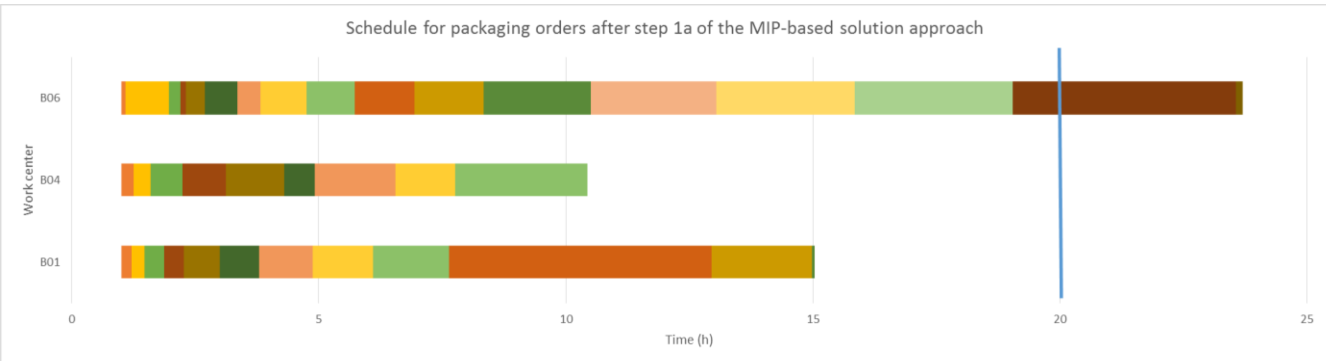

Figure 8.3: Schedule after step 1a of the sequential approach

Figure 8.3 shows the schedule after step 1a of the solution approach. It is clear that the schedule finishes outside of the normal working hours and that it should thus be adjusted to be feasible. Step 1b will have to be performed.

Step 1b

During step 1b, the variables associated with orders that are assigned to the B06 production lines will be relaxed. This means orders assigned to this machine can be freely assigned to the other production lines to reduce the work load on the B06 line.

After unfixing the set of variables, the model is solved again. The solution found during step 1a of the approach is used as a starting point. The main characteristics of the solution that is found after this step is performed are given in Table 8.10 and Figure 8.4.

Table 8.10: Main solution characteristics of the packaging orders after step 1b of the solution approach

Description	*Value*
Used penalties in objective (p_1, p_2 *and* p_3)	(1,1,1)
Best Solution (Objective value)	**347,74**
Best LP bound	**197,31**
% Gap	44.76%
Solution time until feasible (s)	Feasible at start
Total solution time	711s
# of Constraints	5641
# of Variables	1174
# of orders that are scheduled late (Finishing time > D_w)	0
Sum of workload scheduled after shifts have finished	0 h
Total amount of idle time (h)	0
Total changeover penalty	94
# of orders assigned to non-preferred packaging lines	3

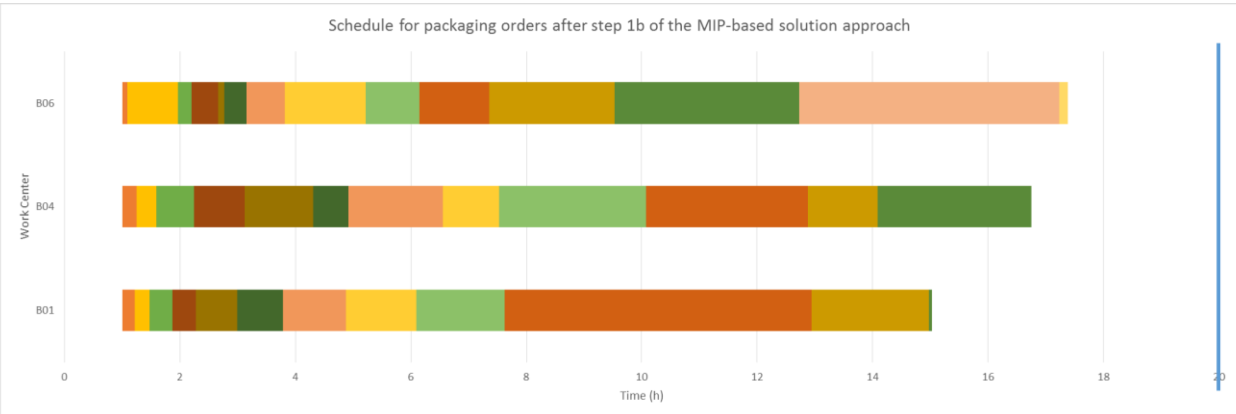

Figure 8.4: Schedule after step 1b of the sequential solution approach

The newfound solution is now feasible because all work centers finish before the end of the last shift ($D_w = 20$). Table 8.10 further shows that only 3 orders were assigned to another line to achieve this result and that the total changeover penalty even decreased.

Step 2

Now that a feasible schedule for the packaging orders has been found, the orders for the intermediate products will be scheduled in step 2. From the schedule of the packaging orders, the times at which the intermediates for these orders should be available can be calculated. These times will serve as the preferred delivery times (D_w) of the pre-processing orders.

The model presented in section 7.2 will be used to solve the problem for the pre-processing orders. Because this approach will be compared to the other approach, the weights in the objective function will remain the same. Later in this section, the effects of using different weights will be shown.

Figure 8.5 and table 8.11 show the results of this step of the approach. The solution shows that the model can quite easily adapt to the schedule of the packaging orders. Only 3 of the orders are scheduled too late for their respective packaging orders.

Table 8.11: Main solution characteristics of the pre-processing orders after step 2 of the solution approach.

Description	*Value*
Used penalties in objective (p_1, p_2 and p_3)	(1,1,1)
Best Solution (Objective value)	160,04
Best LP bound	97,94
% Gap	38.80%
Solution time until feasible (s)	0,55s
Total solution time	730,52s
# of Constraints	1942
# of Variables	545
# of orders that violate the intermediate to final product connection (Starting time $> D_w$)	3
Sum of workload scheduled after shifts have finished	3,15h
Total amount of idle time (h)	0,0024h
Total changeover penalty	55
# of orders assigned to non-preferred packaging lines	0

Figure 8.5: Schedule of all pre-processing orders after step performing step 2 of the sequential approach.

Step 3

The weights that were chosen in step 2 of the approach result in a total of 3 intermediate-final product connection infeasibilities. This means that intermediates are not scheduled to be available on time for use in the final products. At this point a choice will have to be made to either schedule the orders that result in infeasibilities ahead of time, or to adjust the schedule for the packaging orders. For comparative reasons, the only choice presented here is the adjustment of the schedule found in step 1b for the solution that was found in step 2. Scheduling the infeasible orders ahead of time would give the method that is tested here an unfair advantage over the first method. A number of orders would be removed from the problem instance entirely. Step 3 will thus be used to adjust the schedule found in step 1b in order to arrive at a solution that is overall feasible.

The solution that was found in step 1b is used as a starting point for this step. The starting times of the intermediates that were found in step 2 are entered into the parameter G_w. Using this parameter constraints (7.27) now enforce that the connection between the intermediates and final products is feasible.

The model from section 7.3 is solved again with this information. The result is the final solution that is feasible and can be compared to the solution found using the first approach. An overview of the solution is given in table 8.12 and figure 8.6s

Table 8.12: Main solution characteristics of the packaging orders after step 3 of the solution approach

Description	*Value*
Used penalties in objective (p_1, p_2 and p_3)	(1,1,1)
Best Solution (Objective value)	383,05
Best LP bound	294,07
% Gap	23,23%
Solution time until feasible (s)	During pre-solve (<0,58s)
Total solution time	714,07s
# of Constraints	5641
# of Variables	1174
# of orders that are scheduled late (Finishing time > D_w)	0
Sum of workload scheduled after shifts have finished	0 h
Total amount of idle time (h)	2,30 h
Total changeover penalty	102
# of orders assigned to non-preferred packaging lines	3

No major adjustments were made in order to make the full production schedule feasible. Compared to the solution found during step 1b (table 8.10), there was a small increase in the changeover penalty (+8) and a small increase in the idle time on production line B01 (+ 2,30 h).

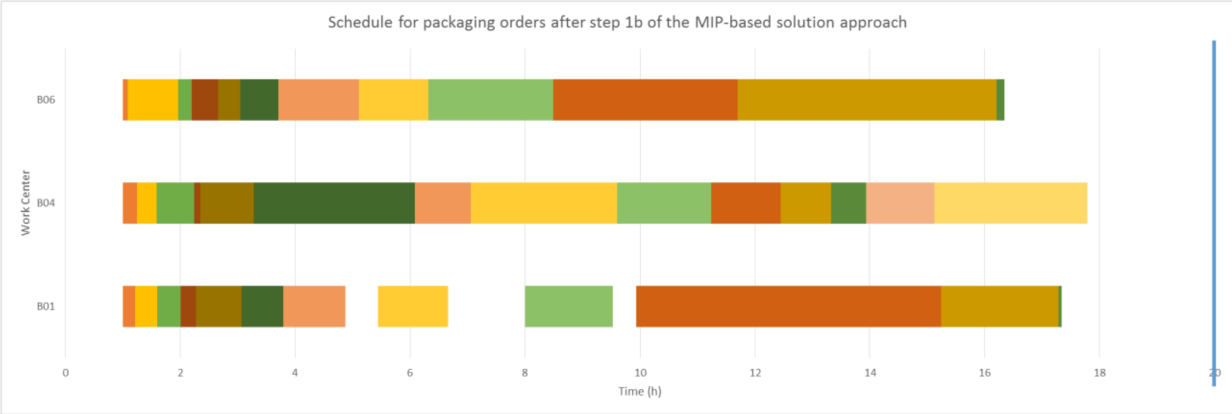

Figure 8.6: Schedule for the packaging orders after step 3 of the sequential approach

Complete solution

In the sequential approach the packaging and pre-processing orders are scheduled separately. In step 2, a schedule was found for all pre-processing orders and in step 3 a final feasible schedule for the packaging orders was found.

These two schedules can now be merged to form the full schedule for the problem instance. This schedule can be compared to the solution found using the sequential approach (section 8.2.1). The result is shown in table 8.13 and figure 8.7.

Table 8.13: Main solution characteristics of the packaging orders after step 3 of the solution approach

Description	*Value*
Used penalties in objective (p_1, p_2 and p_3)	(1,1,1) all problems
Best Solution (Objective value)	583,67
Best LP bound	281,72
% Gap	51,73%
Solution time until feasible (s)	n/a (multiple models)
Total solution time	2869,59 s (all steps)
# of Constraints	n/a (not a single model)
# of Variables	n/a (not a single model)
# of orders that are scheduled late (Finishing time > D_w)	0
Total amount of idle time (h)	2.30 h
Total changeover penalty	157
Changeover penalty for packaging orders	102
Changeover penalty for pre-processing orders	55
# of orders assigned to non-preferred packaging lines	3

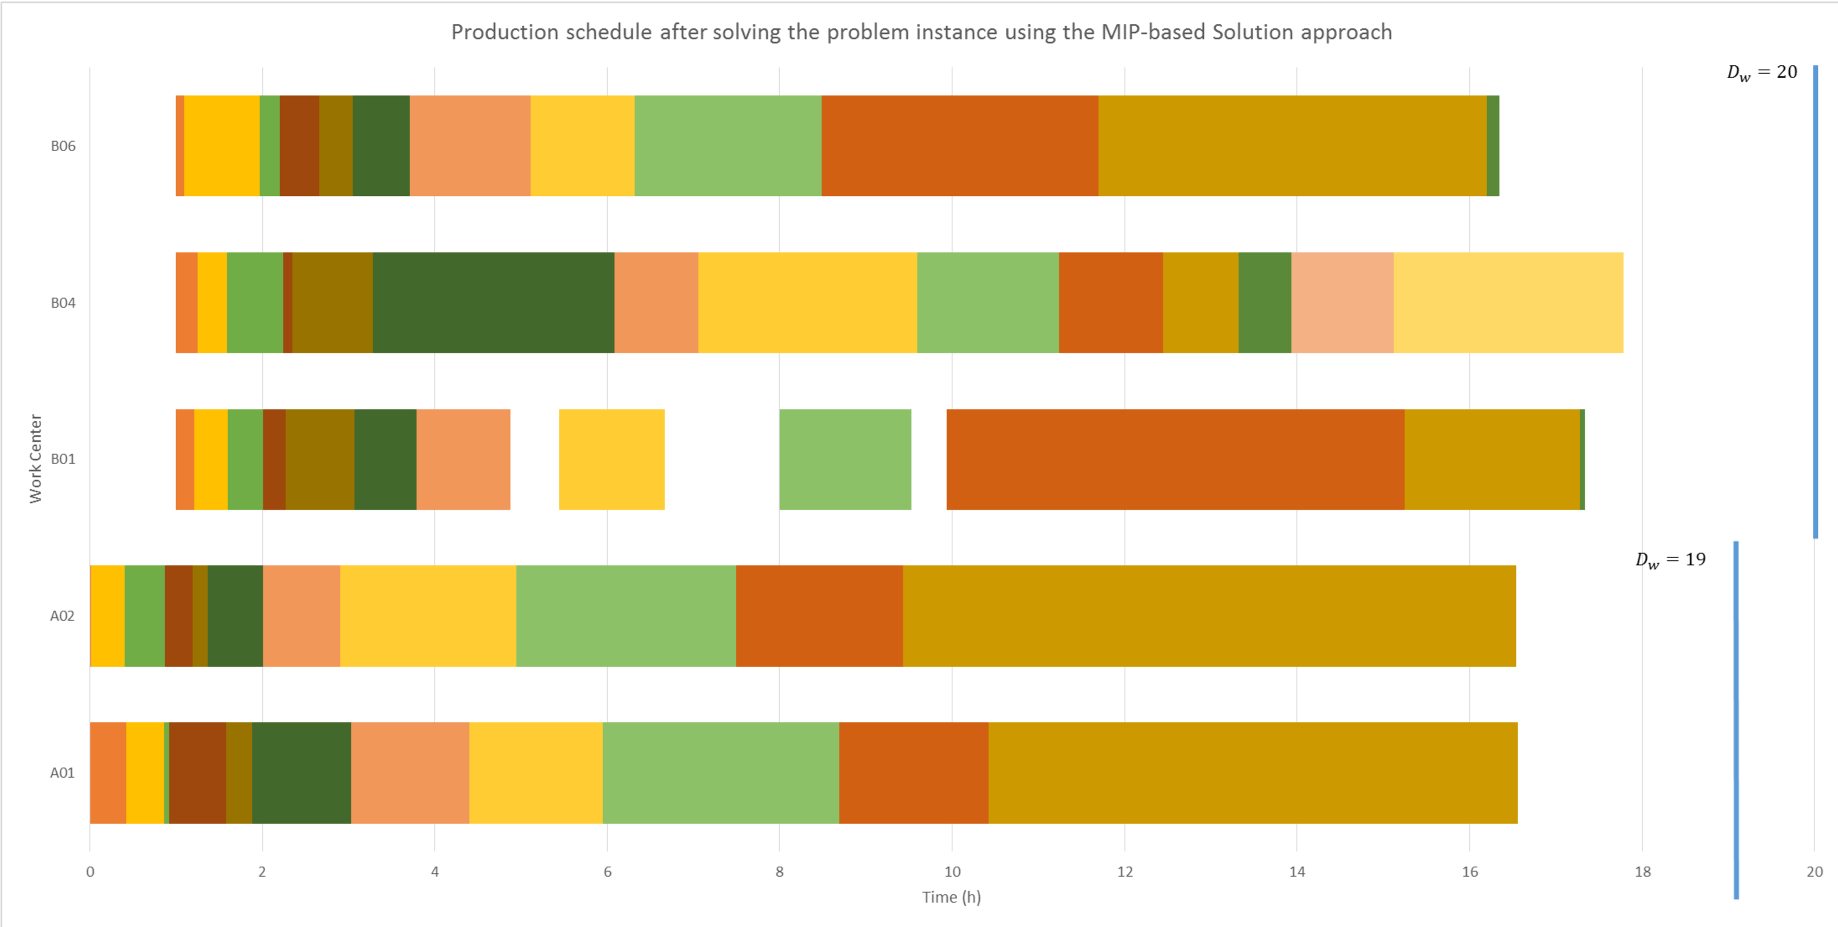

Figure 8.7: Complete production schedule after solving the problem instance using the sequential approach (Section 7)

8.2.3 Comparing the two approaches

Now that solutions have been found for an exemplary problem instance using both approaches, a comparison can be made. The comparison will be made using a few key performance indicators. These indicators are:

1. Total time required to find the presented feasible solution
2. Total time used for solving the model(s)
3. The number of orders that are scheduled to finish after a shift has ended
4. The total changeover penalty
5. Total changeover penalty for the pre-processing orders
6. Total changeover penalty for the packaging orders
7. The amounts of idle time that are part of the schedule (the time during which work centers are not utilized).
8. The amount of orders that are not assigned to their preferred work center.

These KPIs should give a good view of the capabilities of both approaches. Figure 8.8 shows a histogram in which each of the KPIs are compared. All data has been normalized for it to be easily compared. In the histogram all values for the sequential approach form the basis and are equal to 1. This means that if the value for the first approach is < 1, then it is smaller than that value for the MIP-based approach.

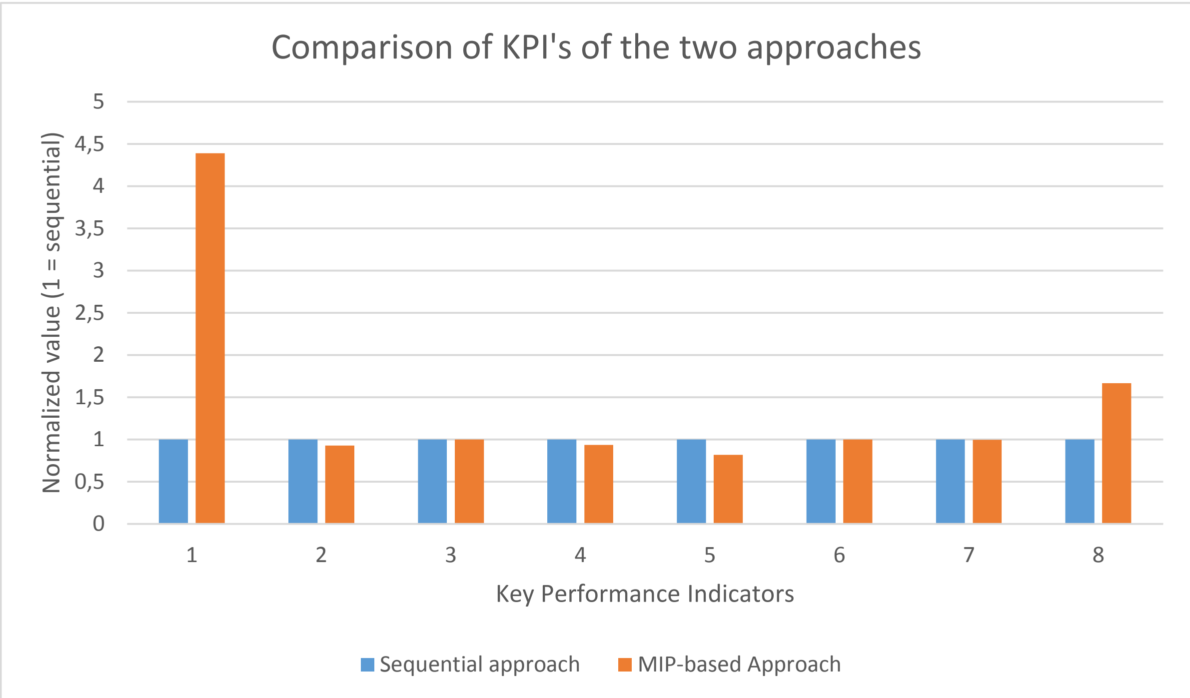

Figure 8.8: Comparison of Key Performance indicators of both solution approaches.

It can be easily seen that the sequential approach performs at least as good as or better than the MIP-based approach for all but some indicators. KPI 2 shows that the total solution times used to solve both models lie close to each other. This is however the result of the choice of stopping criterion. For the sequential approach, almost 84 % of the total solution time is used to check if the stopping criterion has been met (22% for the MIP-based approach). When only the time used until the presented solutions were found is compared (KPI 1), it can be seen that the sequential approach outperforms the MIP-based approach by a large amount. The next section will show that the waiting used in the stopping criteria can be significantly reduced without affecting the solution much. The result of this is that the sequential approach requires less solution time than the MIP-based approach.

Apart from the solution time, the solutions found by both approaches are relatively equal. The only other differences can be found in Key Performance Indicators 4, 5 and 7. This is a direct result of the sequential solving of parts of the problem instance in the sequential approach. In contrast to the MIP-based approach, where orders are all scheduled at the same time, the sequential approach schedules the packaging orders before the pre-processing orders. This means that the packaging orders will have more 'freedom' or scheduling possibilities than the pre-processing orders. The result is a higher changeover penalty for the pre-processing orders (KPI 4 and 5) when the sequential approach is used.

A positive result of this sequential approach is that more packaging orders are assigned to their preferred production lines (KPI 8).

One thing that is not directly clear from figure 8.8 are the differences in changeovers for the packaging department in both solutions. Whereas both solutions have the same total changeover penalty for this department (KPI 6), the individual changeovers are not equal. This can be seen in table 8.14.

Table 8.14: Comparison of the occurrence of different changeover penalties for the packaging departments in the schedules found by the two approaches

Individual changeover penalty	*# of occurrences of this changeover in the MIP-based approach*	*# of occurrence of this changeover in the sequential approach*
1	5	4
2	28	25
4	16	5
5	6	9
6	0	3
10	4	2
Total changeover penalty:	**147**	**157**

As can be seen in the table, both solutions have changeover penalties that lie close to each other. However, the solution found using the sequential approach contains only half the amount of format part changes (which have a penalty of 10). The exchange of format parts on production machines is considered to be the worst possible changeover. Because of this, even though both total changeover penalties are equal, the sequentially generated solution schedule could be perceived as being more convenient.

Overall, the schedules produced by both approaches do not differ much. Both schedules have the characteristics of a good schedule. They both contain little idle time, delays or bad changeovers. However, the sequential approach has a clear advantage over the MIP-based approach. The solution time required to find the solution for the first is a lot less than the time it requires to solve the problem using the latter. Furthermore, the sequential approach contains the possibility to individually steer the solutions for the pre-processing and packaging departments during the different steps. Because the management of the case-company has indicated that the packaging department should be the leading factor in the scheduling process, this can be considered as an advantage.

8.2.4 Progress of the solution while solving the model.

Until now, only final solutions have been presented. However, while solving a model it is interesting to see how quickly the objective function converges to the final solution. When the solution quickly converges, it may be interesting to stop the solver early. This will still result in solutions that lie close to the final solution but would require considerably less time to find.

To see the development of the objective function during the solving of the model, the logger of the solver was used to create a log file that contains all feasible solution and the times at which they were found. From this data a graph can be created that shows the course of the objective function during solving.

Logs were created while solving the comparison problem presented in this section for both the MIP-based approach and step 1a of the sequential approach. Figure 8.9 shows the results for the first.

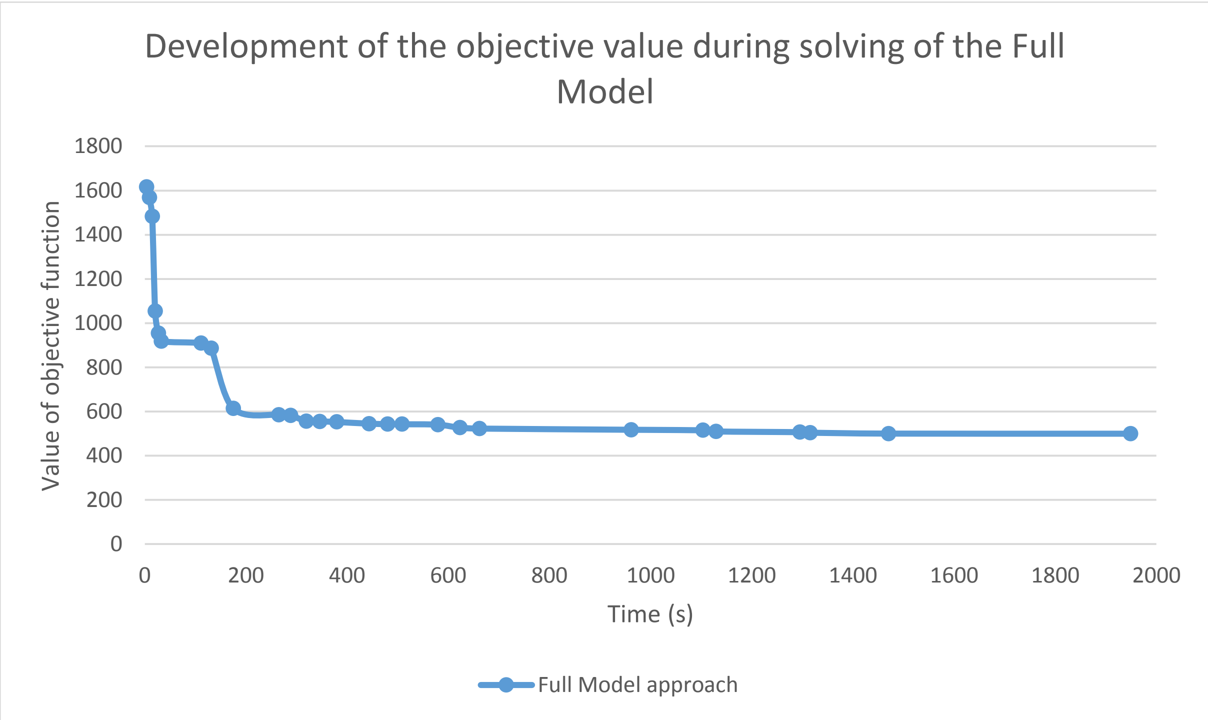

Figure 8.9: Progress of the objective value while using the MIP-based approach to solve the comparison problem.

Figure 8.9 shows that the objective value improves rather rapidly at the start of the solving process. Almost 95% of the improvement is achieved within the first 25% of the time used to arrive at the final solution. This indicates that the solver could have been stopped a lot earlier and the resulting solution would still be close to the final solution.

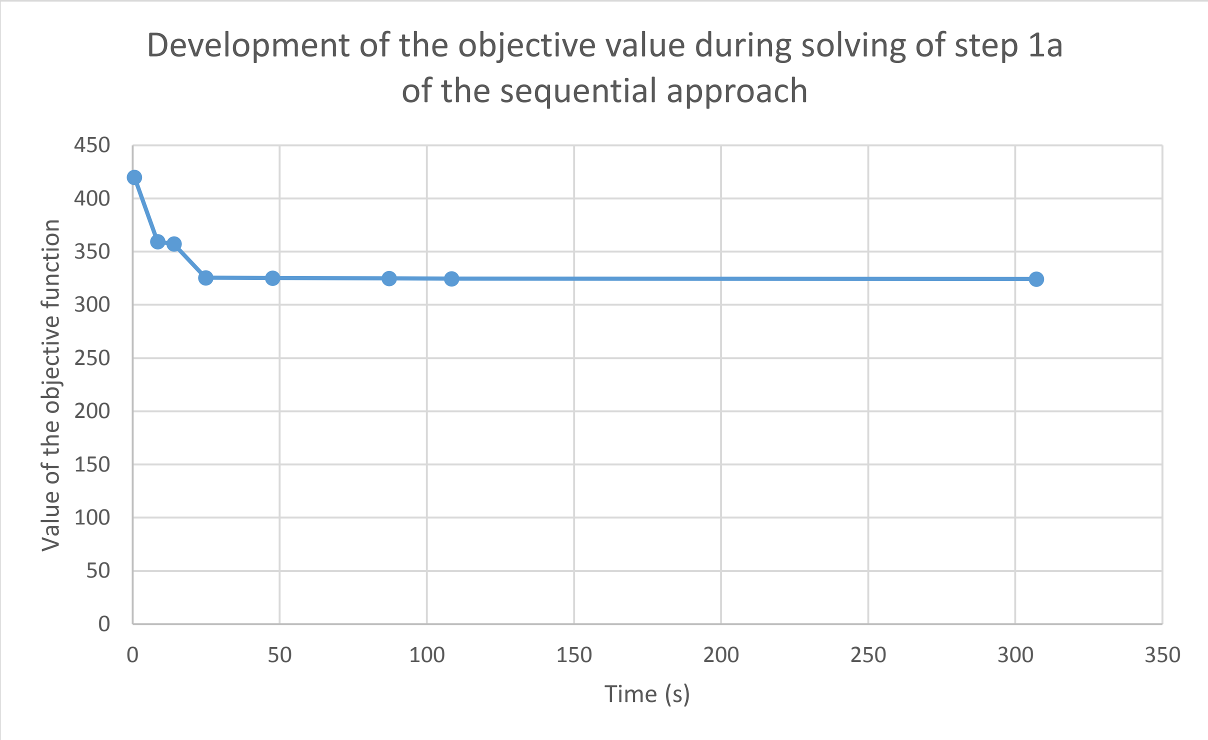

Figure 8.10: Progress of the objective value while using step 1a of the sequential approach to solve the comparison problem.

Figure 8.10 shows the same graph but for the sequential approach. The rapid convergence is even clearer in this graph. Almost 99% of the improvement is made in the first 25 seconds of the procedure. This certainly indicates that using a stopping criterion that stops the solver more quickly, would not degrade the solution.

Results that are similar to the ones presented here were found for the other instances of the problem.

Plenty of time is available to solve the scheduling problem. The required orders are known one day in advance. A solving time of one or two hours to find a good solution would still be reasonable. However the results presented here indicate that using that time may not be necessary when solutions that are just slightly worse are accepted.

This knowledge can be useful when the size of the problem increases or when orders have to be rescheduled in a short amount of time after a change has been made to the orders.

8.3 Solving the real-world industrial-scale problem at Heemskerk

The preliminary results from section 6.2 have already shown that simply solving the model with a commercial solver is not possible for the large problem instances that are encountered at Heemskerk Section 8.2 has shown that the sequential approach that was presented in section 7 performs very well when compared to the MIP-based approach. The solutions that are found using the sequential approach are of similar quality and only require a lot less solution time.

The only viable solution to solve the problem instance of a single production day at Heemskerk is therefore the sequential approach. In this section, this approach will be used to find a production schedule for a representative real production day. The data that is used for the problem instance presented in this section is the actual real-world data from the production process at Heemskerk.

The day that was chosen as a representative production day was a busy day on Wednesday the 11th of July 2013. The Wednesday is usually the 2nd busiest day of the week and should therefore provide an appropriate scheduling challenge.

The main characteristics of production orders that were required to be scheduled on the 11th of July 2013 are presented in table 8.15.

Table 8.15: Main characteristics of the production day of 11-07-2013

Description	*Value*
Total # of orders required to be scheduled	357
# of packaging orders	235
# of bag-packaging orders	146
# of bowl-packaging orders	56
# of high care packaging orders	33
# of pre-processing orders	122
Total number of packaging lines	16
Pre-processing lines	11
Total packaging workload	234,61 h
Available packaging time	304
Utilization of packaging lines	75,93%
Total pre-processing workload	98,46
Available pre-processing time	209
Utilization of pre-processing lines	47,59%

The objective function used in both models of the sequential approach contains three separate parts. As explained before, these three parts represent the following process aspects:

1. The total changeover penalty (weight factor = p_1)
2. The amount of idle time in the production schedule (weight factor = p_2)
3. The of time the schedule finishes after the end of the shift (weight factor = p_3)

In this section, weights will be chosen that generally give good results during the solving of different instances of the problem. The next section will show the actual influence of these weights on the solutions found by the models.

To measure the quality of found production schedule, a number of KPIs will be given for the solution after each step of the solution process. These KPIs are:

1. Total changeover penalty score
2. Total amount of idle time in the schedule in hours
3. Total number of orders that are delivered too late. These orders finish either after the shift has ended (step 1a, 1b and 3) or finish too late to be used in their respective packaging orders (step 2, pre-processing orders).
4. Sum of workload that is scheduled after normal operation hours (total late time).
5. Number of packaging orders that are not assigned to their preferred work center.

In this section, the 4 individual steps (1a, 1b, 2 and 3) of the sequential approach will be discussed in succession. For each of the steps, the resulting solution will be presented along with a visual representation of the schedule that is found. This will give a good example of the progress that is made in each step of the approach. Finally, the final production schedule that is found for the complete problem instance will be presented.

8.3.1 Step 1a

In step 1a all packaging orders are scheduled. During this step, all orders are assigned to their preferred work centers. This is done by fixing the assignment variables. A more thorough explanation of this step can be found in section 7.

Because there is no connection between the 3 different packaging departments, the packaging orders can be split into three parts. Table 8.15 shows the size of the problem instance for each of the three departments. Each of these three smaller problem instances is solved separately using the CPLEX solver. Afterwards, the solution for all packaging orders is found by recombining the results for each of the three instances.

The first and largest packaging instance includes al orders of products that are packed into bags. A total of 146 orders have to be scheduled on 11 work centers. The second instance includes all 56 orders that are packaged into bowls. These orders have 3 work centers available to them. The last instance contains 33 HighCare work orders that need to be assigned to two different machines. The results for these three separate instances are shown in table 8.16. Each column represents one of the instances.

Table 8.16: Main solution characteristics for the three instances of packaging orders after step 1a of the solution approach

Description	*Bag-packaging orders*	*Bowl-packaging orders*	*HighCare packaging orders*
Used penalties in objective (p_1, p_2 *and* p_3)	(1,0.01,1)	(1, 0.01, 1)	(1, 0.01, 1)
Best Solution (Objective value)	460,64	121,97	66,67
Best LP bound	358,98	105,40	56,74
% Gap (%)	22,07	13,58	14,91
Solution time until feasible (s)	During pre-solve (<6,53)	During pre-solve (<0,99)	During pre-solve (<1.39)
Time until solution was found (s)	134,53s	36,72s	29,53
Total solution time (s)	734,53	636,72	629,53
# of Constraints	65515	7707	4475
# of Variables	2989	1241	1321

The solver was stopped using the same stopping criterion as in section 8.2. This was done in order to present the quality of the solutions that can be achieved with the presented approach. As explained before, the waiting time for the stopping criterion can be reduced without affecting the solution much. All presented total solution times in this section include the 600 seconds during which no change in the best found feasible solution occurs.

From the results in table 8.16 it is clear that the bag-packaging instance is the largest and most difficult of the three instances to solve. The other two instances are considerably smaller in size and take less time to solve. Nevertheless, the solution times are very acceptable.

It should be noted that because HighCare packaging orders have no preferred production lines (both available lines are identical). Therefore the assignment variables were not fixed. This also explains why the smaller HighCare instance contains more variables than the bowl-packaging instance. Because only 33 orders need to be scheduled for this instance, this does not result a large increase in solution times.

To show the combined solution for all packaging orders that was found during this step, the solutions from all three instances are combined. Table 8.17 lists the resulting solution characteristics.

Table 8.17: Main solution characteristics for the solution of all packaging orders after step 1a of the solution approach.

Description	*Value*
# of orders that are scheduled late (Finishing time > D_w)	12
# of work centers that finish after the last shift ends	3
Sum of workload scheduled after shifts have finished	18,03
Total amount of idle time (h)	0
Total changeover penalty	566
# of orders assigned to non-preferred packaging lines	0
Time until presented solution was found (sequential)	200,78 s
Total solution time (sequential)	2000,78 s
Total solution time (simultaneously)	734,53 s

The results show that the actual solution is infeasible at this point. 12 production orders on 3 different work centers are scheduled after the last production shift has ended. During this step, these orders cannot be moved because the assignment variables are fixed. All work orders are also assigned to their preferred production line because of this.

Figure 8.11 gives a visual representation of the production schedule that was found. Each colored bar again indicates an individual work order assigned to that machine in that timeslot.

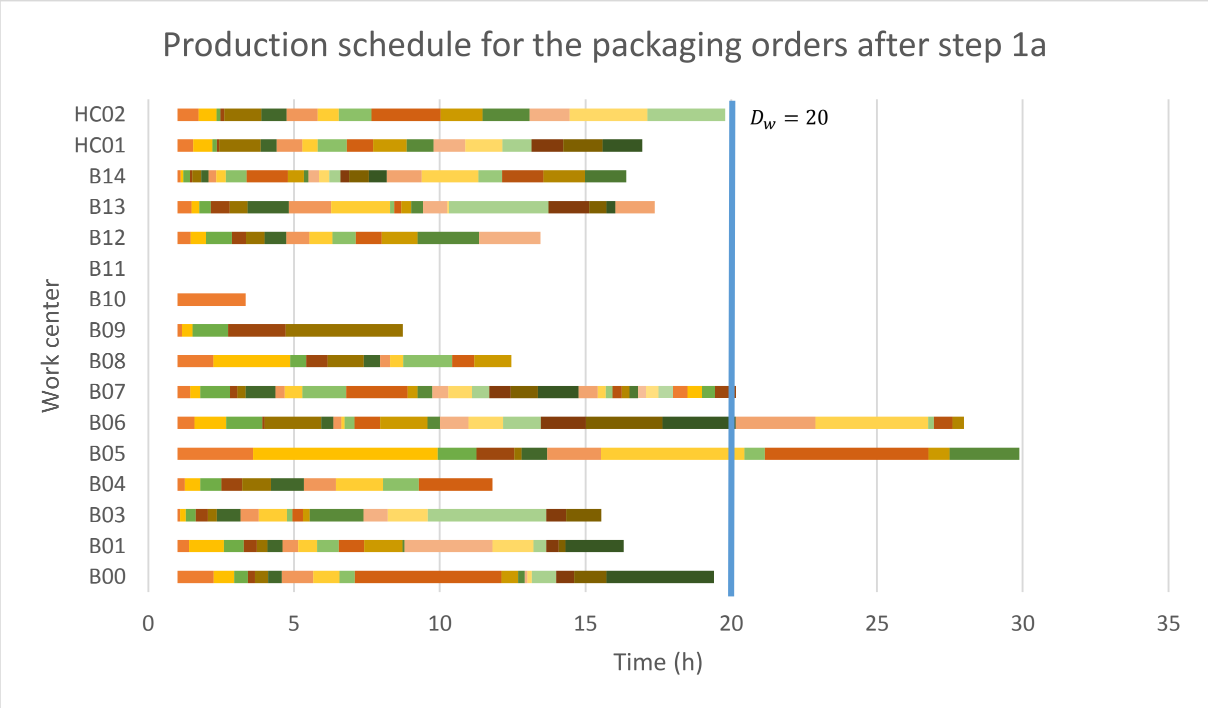

Figure 8.11: Production schedule for all packaging orders after step 1a of the solution approach.

The figure clearly shows that the B07, B06 and B05 production lines finish after the last shift ends ($D_w = 20$). The B11 production line has no orders assigned to it. This is because no orders have the B11 as their preferred production line. One could call this production line a 'Jack of all trades, master of none'.

8.3.2 Step 1b

It became clear in step 1a that the solution for the packaging orders was still infeasible. During step 1b, the orders that are assigned to late finishing work centers will be unfixed. This means that the assignments variables of all orders that are assigned to these three machines are unfixed. The orders on these machines are thus allowed to freely be assigned to other work centers on which they can be processed.

Because all work centers that finish late are part of the bag-packaging department, only the solution for this instance will have to be adjusted in this step. The solutions for the other two departments can be reused.

The solution found in step 1a is used as a starting point in this step. After all required variables are unfixed, the model is solved again using the CPLEX solver. The results after solving this instance are shown in table 8.18

Table 8.18: Main solution characteristics for the bag-packaging instance after step 1b of the solution approach

Description	*Bag-packaging orders*
Used penalties in objective (p_1, p_2 and p_3)	(1,0.01,5)
Best Solution (Objective value)	362,63
Best LP bound	286,01
% Gap (%)	21,13
Solution time until feasible (s)	feasible at start
Time until presented solution was found (s)	1138,63
Total solution time (s)	1738,63
# of Constraints	137252
# of Variables	9599

This table shows that both the amount of variables and the amount of constraints increases dramatically by unfixing the assignment variables on 3 machines. This considerably increases the solution time required for this step. Unfixing the variables on each of the 3 overburdened machines sequentially could reduce the solution time required. In this case, first all orders on the most overburdened production line would be relaxed. After resolving the model, all orders would be fixed again except for the orders on the 2nd most overburdened machine. This will decrease the total solution time required. For now, this is unnecessary because 20 minutes is still a reasonable solving time for this part of the problem.

Table 8.19 lists the performance indicators of the full solution after the orders of the bag-packaging department have been adjusted in this step.

Table 8.19: Key Performance Indicators for the solution of all packaging orders after step 1b of the solution approach.

Description	*Value*
# of orders that are scheduled late (Finishing time > D_w)	0
# of work centers that finish after the last shift ends	0
Sum of workload scheduled after shifts have finished	0
Total amount of idle time (h)	0
Total changeover penalty	534
# of orders assigned to non-preferred packaging lines	28 (of 235)
Time until presented solution was found (s)	1138,63
Total solution time (s)	1738,63

The table shows that all work that was scheduled outside of regular work times has been rescheduled to other machines. For this 28 orders were moved to different production lines. This has also resulted in a reduction of the total changeover penalty. This is mainly the result of the usage of the B11 production line. This line did not have any orders assigned to it before, when the number of used production lines increases, the number of changeover automatically decreases as well.

Figure 8.12 shows the schedule of all packaging orders after they have been rescheduled in this step. It is clear that now $D_w < 20$ for all work orders.

Figure 8.12: Production schedule for all packaging orders after step 1b of the solution approach.

8.3.3 Step 2

In step 2 all pre-processing orders are scheduled. This is a relatively straightforward process. All orders for this department are scheduled at once. The main characteristics of the solution for this step are given in table 8.20

Table 8.20: Main solution characteristics for the solution of all packaging orders after step 2 of the solution approach.

Description	*Value*
Used penalties in objective (p_1, p_2 and p_3)	(1, 0.01, 10)
Best Solution (Objective value)	616,61
Best LP bound	387,39
% Gap (%)	37,17
Solution time until feasible (s)	7.23
Time until presented solution was found (s)	902,04
Total solution time (s)	1502,04
# of Constraints	31860
# of Variables	2187
# of orders that are scheduled late for usage in their respective final production(Finishing time > D_w)	7
# of work centers that finish after the last shift ends	0
Sum of workload scheduled late	11,43
Total amount of idle time (h)	2,97
Total changeover penalty	424

The table shows that only 7 of the 122 pre-processing orders are scheduled too late for usage in their respective final products. Furthermore there is relatively little idle time (2,97h) on the production lines.

Figure 8.13 shows the schedule of the pre-processing orders after they have been scheduled in this step.

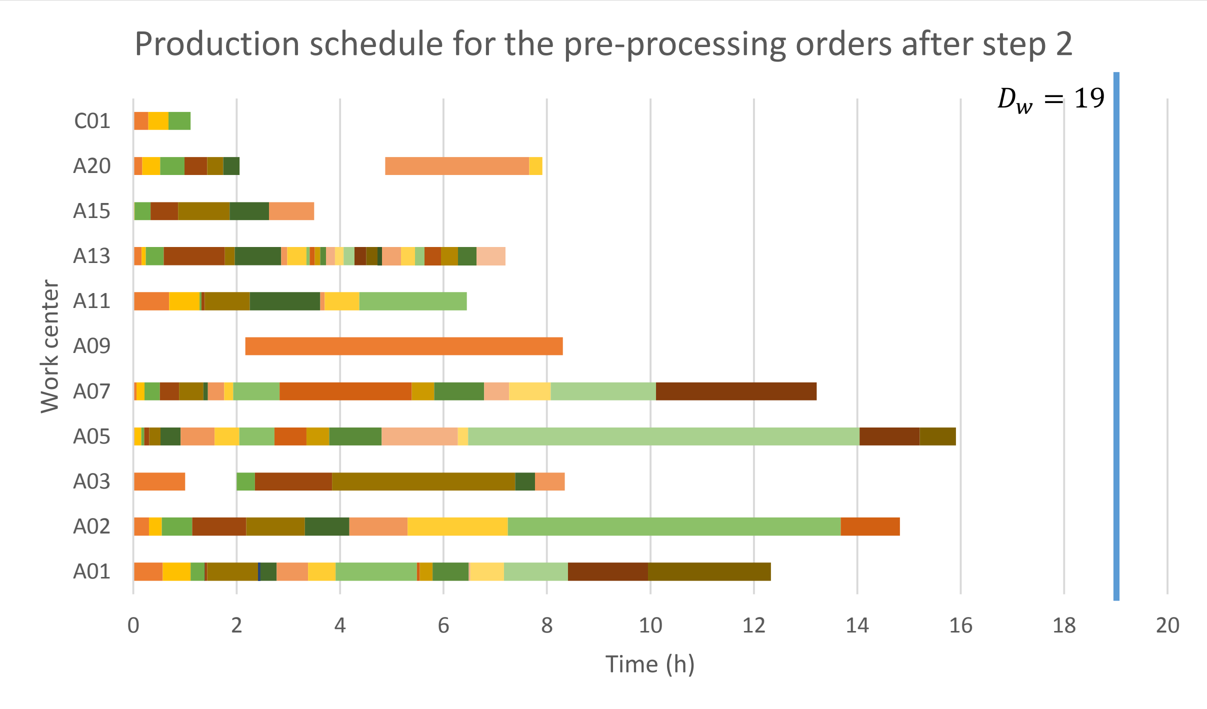

Figure 8.13: Production schedule for all pre-processing orders after step 2 of the solution approach.

8.3.4 Step 3

In this step the complete schedule will be adjusted to be overall feasible. The schedule found in step 2 of the approach results in 7 intermediates that are not scheduled on time for usage in the final products they are required in.

At this point, the choice could be made to reschedule those 7 orders as 'pre-work' orders at the end of the previous production day. These orders would thus be brought forward by one day. After doing this, both the schedule for the pre-processing orders found in step 2 and the schedule for the packaging orders found in step 1b are feasible and would need no further adjustments.

As was presented in section 7, bringing the infeasible pre-processing orders forward is often a good choice that has few negative side effects. However, to show the capabilities of the approach, it will be assumed that these orders cannot be brought forward. The result of this is that the solution of the packaging orders needs to be readjusted for late delivery of those 7 pre-processing orders.

To adjust the schedule, the parameter G_w will contain the starting times of all the pre-processing orders found in step 2. The model for the packaging orders (section 7.3) will then be re-solved. Constraints (7.27) then enforce that packaging orders can only be scheduled after their required intermediates are available.

Table 8.21 lists the main characteristics of the model after it was re-solved. Because the infeasibilities only occur in the bag-packaging department, only this instance has to be solved again. The solution that was found in step 1b is used as a starting point for this solution.

Table 8.21: Main solution characteristics for the bag-packaging instance after step 3 of the solution approach.

Description	*Bag-packaging orders*
Used penalties in objective (p_1, p_2 and p_3)	(1,0.01,5)
Best Solution (Objective value)	364,38 (+1,75)
Best LP bound	287,99
% Gap (%)	20,96
Solution time until feasible (s)	Feasible at start
Time until presented solution was found (s)	1010,89
Total solution time (s)	1610,89
# of Constraints	137252
# of Variables	9599
# of orders that are scheduled late (Finishing time > D_w)	0
# of work centers that finish after the last shift ends	0
Sum of workload scheduled after shifts have finished	0
Total amount of idle time (h)	7,28 h (+7,28)
Total changeover penalty	534 (+/- 0)
# of orders assigned to non-preferred packaging lines	42 (+14)

As can be seen from the table, the impact of adjusting the schedule is significant. An additional 14 orders were assigned to non-preferred production lines and a total of 7,28 hours of idle time is introduced. Figure 8.14 shows the schedule that was found.

The schedule shows that a lot of idle time is introduced on production line B09, other lines see much less or no idle time at all. Generally, idle time on packaging lines is unwanted. However, the small amount of idle time that is introduced to, for example, machine B03 may not be much of a problem. Ultimately, it is a manual decision on how to treat these infeasibilities. The company will have to create a set of rules for this. They may want to move forward as many pre-processing orders as needed or they may only want to move the orders that result in a lot of idle time (B09). This is something that will need to be decided by the management.

Overall, as will be shown in the next section, the re-adjustment of the packaging orders in step 3 will result in quite some idle time being introduced to the schedule.

Figure 8.14: Production schedule for all packaging orders after step 3 of the solution approach

8.3.5 Complete solution

The solution for the packaging orders found in step 3 and the solution for the pre-processing orders found in step 2 can be combined to form the complete feasible production schedule for the selected production day. This solution is the combination of all steps of the sequential approach. The total solution time is the sum of all steps of the solution. All other performance indicators can be summed as well.

Table 8.22 shows the full list of results for the complete solution. On the next page, the schedule that is found for all production lines is shown in figure 8.15.

Table 8.22: Main solution characteristics for the bag-packaging instance after step 3 of the solution approach.

Description	*Bag-packaging orders*
Sum of workload scheduled after shifts have finished	0
# of orders that finish outside regular production times	0
Total amount of idle time (h)	10,25
Total changeover penalty	958
# of orders assigned to non-preferred packaging lines	42
Time until presented solution was found (s)	3251,34
Total solution time (s)	5651,34

The table shows that the total required solution time to find the complete schedule is less than one hour. Due to the generous stopping criteria a little over 93 minutes were used solving the problem. Since the definitive workload for the next day is known circa 12 hours in advance, the problem can be solved in a reasonable amount of time.

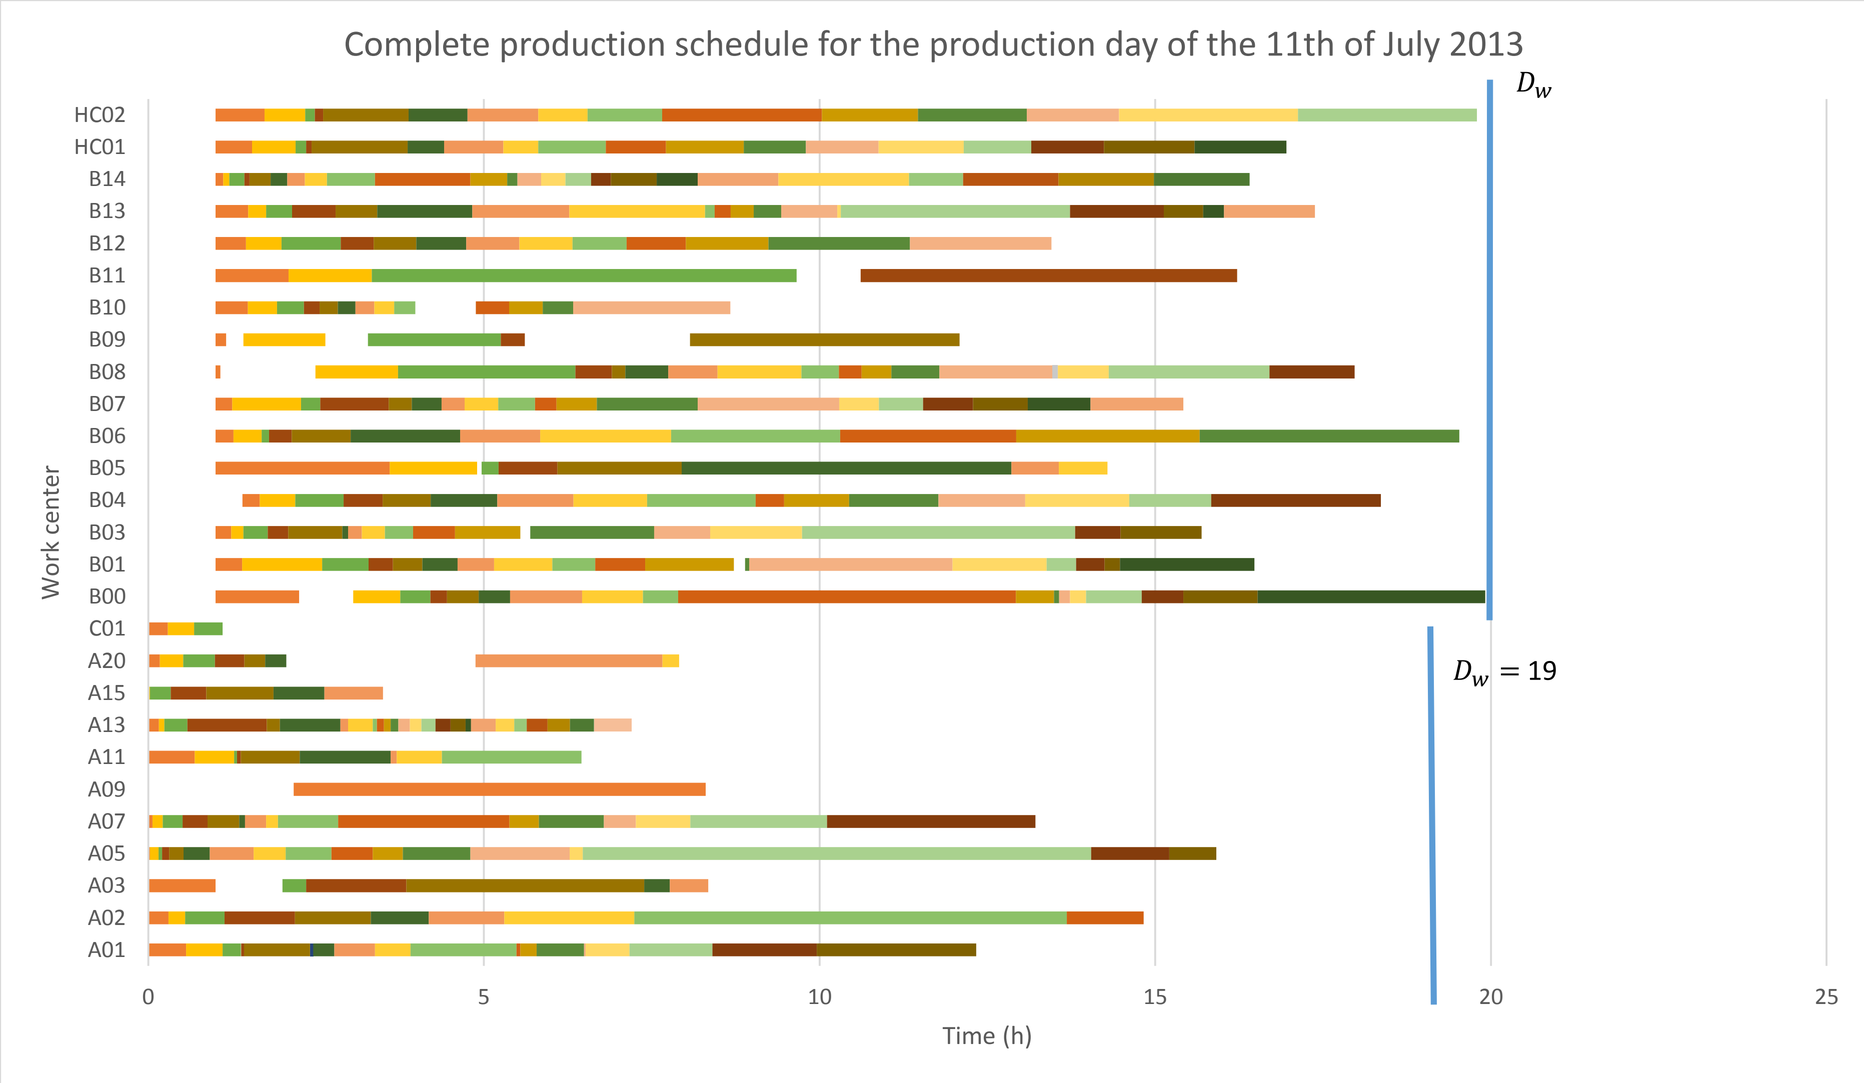

Figure 8.15: Complete schedule for the production day of 11-7-2013.

8.3.6 Results for different production days

In order to confirm that the sequential approach gives similar results when used on the data of other production days, the problem instances of two additional days were solved. These days are the 10th of July 2013 and the 10th of October 2013.

For both these instances, no significant change in either the solution times or the quality of the found solutions was found. This suggests that it should be possible to use the presented approach for the data of any production day.

8.4 Influence of the penalty factors in the objective function

The objective function of the mathematical model contains three individual parts. Each of these parts has a weight factor associated with it. These weight factors are used to steer the solution of the model. Ideally, when one of the weight factors is higher than the others, the aspect associated with that factor should be more profound in the solution that is found.

To see the effects of varying these weights in the objective function, a large number of tests has been performed. All these tests were conducted on the data for a full production day at Heemskerk The day used for the tests presented here is the 11th of July 2013. This is the same representative day as was used in the previous section.

For a good review of the influence of the penalty factors in the objective function, all 3 (4 when including 1a and 1b separately) steps of the sequential approach were solved with a range of different weights. For each step, a table with different weights and the main characteristics will be given.

Step 2 and 3 of the solution approach require data input from the previous step of the approach. For the tests in this section, a representative solution that shows good results will be used for the next step when this data input is required. In most cases this data will be equal to the solution found for the problem in section 8.3.

The effect of the weight will be measured by listing the first four of the five key performance indicators that were used to measure the quality of the solution in the previous section. These KPIs are:

1. Total changeover penalty score
2. Total amount of idle time in the schedule in hours
3. Total number of orders that are delivered too late. These orders finish either after the shift has ended (step 1a, 1b and 3) or finish too late to be used in their respective packaging orders (step 2, pre-processing orders).
4. Sum of workload that is scheduled after normal operation hours (total late time).

8.4.1 Step 1a

Step 1a is used to schedule all packaging orders for the first time. In this step all packaging orders have a fixed assignment to their preferred production line. The only unfixed variables are the changeover variable (Y_{wv}) and the starting times (S_w).

Table 8.23 lists all five characteristics of the solutions that were found for this step by varying the different weights. For each different combination of weight factors in this table, the solver was run for 1200 seconds. In all cases presented the objective value of the best-found feasible solution did not change by more than 1 during the last 600 seconds of the solving period.

Table 8.23: Key performance indicators for different weight factors used in the objective function during step 1a of the sequential approach.

Instance no.	(p_1, p_2, p_3)	*Total changeover penalty*	*Total amount of idle time*	*Number of late work centers*	*Sum of all work load scheduled late*
1	**(1, 0.01, 1)**	**566**	**0**	**3**	**18,03**
2	(1, 0.01, 2)	566	0	3	18,03
3	(1, 0.01, 10)	568	0	3	18,03
4	(1, 0.01, 100)	574	0	3	18,03
5	(2, 0.01, 1)	566	0	3	18,03
6	(10, 0.01, 1)	566	0	4	33,23
7	(100, 0.01, 1)	566	0	4	33,23
8	(1, 1, 1)	628	0	3	18,03
9	(1, 0, 0)	566	13817	10	6900

In almost all of the instances, weight factor p_2 was chosen very small. This factor can be this small because during step 1a, the connection with the pre-processing orders is not yet made. This means orders can be freely scheduled in time without being restricted by the availability of their intermediates. This mechanism is explained in figure 6.2 in section 6. Choosing any weight larger than zero already results in all orders being scheduled adjacent in time, i.e. each next order starts right after the previous order has finished. By keeping p_2 this low, the other parts of the objective function have a large influence. Comparing instance 1 and instance 8 can easily show this. The only difference between these instances is the factor p_2. Both solutions have zero idle time and the same amount of late orders. However, the total changeover penalty of instance 8 is 62 points higher than the penalty of instance 1. This shows that by setting p_2 very low, the influence of the other two parts of the objective function can be increased.

From instance 1 to 4, the weight associated with the late delivery of orders is gradually increased. This has no positive effect on the actual late finishing times of orders. This is a direct result of the method used during step 1a. To reduce the number of orders that finish late, the orders have to be assigned to a different machine. Because all orders are fixed to their preferred machine, this is not possible. Therefore p_3 has little influence in this step. Instances 1 to 4 do however show that when $p_3 \gg p_1$, the influence of p_1 is reduced and the total changeover penalty increases.

Instance 5,6 and 7 show a gradual increase of factor p_1 compared to the other factors. It can be seen that this does not result in a lower changeover penalty. Again, when $p_1 \gg p_3$, it can be seen that the influence of p_3 decreases and that more orders are scheduled late.

To find out whether it is possible to lower the changeover penalty below 566, instance 9 was ran. This shows that even when only the changeover penalty is of concern, it will still not go lower than a total of 566. This also indicates that it is no use to use a lower weight for p_3 than the weight used in instance 1.

Overall, instance 1 with the weights (1, 0.01, 1) gives the best results in this step. These weights give the lowest idle time, the lowest number of late orders and the lowest total changeover penalty.

8.4.2 Step 1b

During step 1b the orders that are scheduled on machines that finish late are relaxed. This means the orders on these production lines can be freely assigned to other work centers on which they are allowed to be produced.

Table 8.24 lists the various results that were found using different weights in the objective function during this step

Table 8.24: Key performance indicators for different weight factors used in the objective function during step 1b of the sequential approach.

Instance no.	(p_1, p_2, p_3)	*Total changeover penalty*	*Total amount of idle time*	*Number of late work centers*	*Sum of all work load scheduled late*
1	(1, 0.01, 1)	546	0	3	7,18
2	**(1, 0.01, 2)**	**538**	**0**	0	0
3	**(1, 0.01, 5)**	**538**	**0**	**0**	**0**
4	(1, 0.01, 10)	548	0	0	0
5	(1, 0.01, 100)	550	0	0	0
6	(5, 0.01, 1)	534	0	3	3,62
7	(10, 0.01, 1)	558	0	3	22,4

For the same reasons as during step 1a, weight factor p_2 has been set very low. From instance 1 to 5 weight factor p_3 is gradually increased. It can be seen that only after p_3 has been increased to 2, it has enough influence to complete eliminate all orders that finish late. Increasing this factor any higher has no further influence until the point is reached after which the total changeover penalty suffers from the difference ($p_3 > 5$).

When more weight is given to the changeover penalty (instance 6 & 7), not much is gained. Although a small decrease in changeover penalty is noticed in instance 6, this directly introduces orders that finish late again. This is unfeasible and therefore of no use.
The best results are obtained for instance 2 and 3. Both of these give the exact same results. Because it is very important that orders do not finish late, the solution found by using the weights of instance 3 is used for step 2 of the approach.

8.4.3 Step 2

In step 2 all pre-processing orders are scheduled. The schedule that was found for the packaging orders in step 1b is used to determine the preferred starting time of the orders. The on-time delivery of the pre-processing orders is then maximized by part 3 of the objective function. Increasing the weight of p_3 should thus result in less intermediate to final product infeasibilities.

In contrast to what was seen in step 1a and 1b, weight factor p_2 cannot be set arbitrarily low without influencing the results. Because some intermediates are dependent on other intermediates, not all intermediate orders can be scheduled anywhere in time. An order that requires another intermediate should always be scheduled after that intermediate has finished production. Because of this, weight factor p_2 should directly influence the amount of idle time the production schedule contains.

The characteristics of the solutions that were found using different objective function weights during step 2 are found in table 8.25.

Table 8.25: Key performance indicators for different weight factors used in the objective function during step 1b of the sequential approach.

Instance no.	(p_1, p_2, p_3)	*Total changeover penalty*	*Total amount of idle time*	*Number of orders delivered late*	*Sum of all work load scheduled late*
1	(1, 1, 1)	409	2,56	17	31,47
2	(1, 1, 2)	427	1,2	14	26,58
3	(1, 1, 10)	449	1,81	10	14,32
4	(1, 1, 100)	450	1,18	7	9,86
5	(1, 1, 0)	405	2,56	22	51,16
6	(0, 0, 1)	601	12,5	4	8,58
7	(2, 1, 1)	391	1,97	18	42,45
8	(10, 1, 1)	374	0,96	24	62,78
9	(100, 1, 1)	369	2,02	25	66,35
10	(1, 0.01, 1)	386	1,83	13	22,99
11	(1, 0.01, 2)	395	2,4	10	17,06
12	**(1, 0.01, 10)**	**424**	**2,97**	**7**	**11,43**
13	**(1, 0.01, 100)**	**447**	**2,67**	**4**	**9,79**
14	(2, 0.01, 1)	378	5,04	12	31,23
15	(10, 0.01, 1)	369	3,02	16	59,13
16	(100, 0.01, 1)	369	3,02	16	57,93

The table shows that the influence of p_2 is limited. The amount of idle time varies between 1,18 hours and 6,03 hours when $p_2 > 0$. Neither of these values is very high when one considers that 122 pre-processing orders were scheduled. This is especially so because idle time is not a big concern in this department.

It can also be seen in the table that weight factor p_1 does have a big influence on the total changeover penalty of the solution. Especially for instance 14, 15 and 16, the changeover penalty is reduced significantly when compared to the other instances. This does however result in a lot of orders being delivered too late for their respective final product. To cope with this, a lot of changes to the schedule for the packaging orders would have to be made.

Because the packaging department is leading, having a somewhat increased changeover penalty is acceptable in the pre-processing department. Therefore the solutions found for instance 12 and 13 are considered sufficient. The increase of weight factor p_3 results in a low amount of orders that are scheduled too late. This makes it easy to either produce these orders on the previous day or to adjust the schedule for the packaging orders in step 3.

8.4.4 Step 3

In step 3, a choice is made to either advance the pre-processing orders resulting in infeasibilities to the previous day or to adjust the schedule of the packaging orders to be feasible.

When looking at the influence of the weights, only the last choice is interesting. If orders are advanced to the previous day, the schedule of the pre-processing orders would not have to be rigorously adjusted. Orders that are advanced are simply removed from the problem instance and all other orders can just be moved forward as much as possible. The schedule of the packaging would not have to be adjusted at all. Therefore only the second option, i.e. the adjustment of the scheduling of the packaging orders until they comply with the delivery of their intermediates, will be considered in this step.

To re-solve the model of the packaging orders in this step, input data is required. The starting times of the intermediates are used as input in parameter G_w of the model. Constraints (7.27) then ensure that packaging orders can only be scheduled after their required intermediates are available. Both instance number 12 and 13 from table 8.18 are used as input in this step. These instances give relatively good results with few intermediate – final product infeasibilities and should therefore require fewer adjustments in the schedule of the packaging orders.

The results that were found for the different sets of weight factors are given in table 8.26.

Table 8.26: Key performance indicators for different weight factors used in the objective function during step 3 of the sequential approach.

Instance no.	(p_1, p_2, p_3)	*Used step 2 instance*	*Total changeover penalty*	*Total amount of idle time*	*Number of late work centers*	*Sum of all work load scheduled late*
1	(1, 0,01, 5)	12	536	18,41	0	0
2	(1, 1, 5)	12	614	18,41	0	0
3	(1, 5, 5)	12	690	16,14	0	0
4	(1, 10, 5)	12	792	8,56	2	12,39
5	(1, 100, 5)	12	866	17,81	2	19,03
6	(1, 0.01, 5)	13	542	16,8	2	0,22
7	(1, 1, 5)	13	602	4,05	0	0
8	(1, 5, 5)	13	766	10,8	2	6,8
9	(1, 10, 5)	13	766	12,76	2	8,4
10	(1, 100, 5)	13	872	13,59	0	0

The results show that it is hard to bring the amount of idle time down by raising factor p_2. Although the amount of idle time is somewhat lower for instance number 4, it is higher again for instance number 5. Factor p_2 is therefore not very effective at lowering the amount of idle time in the schedule.

When factor p_2 is higher than factors p_1 and or p_3, the quality of the aspects associated with these last two weights decreases. The total changeover penalty sees a very large increase when the weight factor p_1 is small compared to the other weights. There is no single one solution that stands out from the rest. All instances show a significant amount of idle time. Although instances 2 and 7 contain no orders that finish late, the increase in the changeover penalty from respectively instances 1 and 6 is considerable. At this point it is unclear whether solutions that are significantly better even exist. The comparison in section 8.2 suggests that when the MIP-based approach could be used, results would have been similar.

8.5 Using the feasible solution from the sequential approach in the MIP-based approach

In section 6.4 it was shown that it was not possible to find a feasible solution for presented problem instance in a reasonable amount of time. After more than 13 hours the solver ran out of memory without finding a solution. In the previous section, a feasible solution was found using the sequential approach. It is therefore interesting to see how the initial model from section 6 performs when this feasible solution is used as a starting point for solving the model.

The problem instance of the 11th of July 2013 was reused for this purpose. As a starting point for the solving process, the solution that was found using the sequential approach in section 8.3 was used as the advanced basis for the solver. Table 8.27 shows the results after running the CPLEX solver for 10 hours. The numbers between brackets show the difference with the solution that was found in section 8.3.

Table 8.27: Main characteristics of the problem instance and solution after using a starting point to initialize the MIP-based approach.

Description	*Bag-packaging orders*
Used penalties in objective (p_1, p_2 and p_3)	(1,1,10)
Solution time until feasible (s)	Feasible at start
Time until presented solution was found (s)	34536
Total solution time (s)	36000
# of Constraints	496734
# of Variables	88352
# of orders that are scheduled late (Finishing time > D_w)	2 (+2)
# of work centers that finish after the last shift ends	1 (+1)
Sum of workload scheduled after shifts have finished	2,62h (+2,62)
Total amount of idle time (h)	28,43h (+18,18)
Total changeover penalty	1010 (+52)

The test immediately shows that this is not a viable approach. After running the solver for the large amount of 10 hours, only a small number (<10) of small improvements in the best-found objective value have occurred. Furthermore, the solution is actually worse than the solution that was found using the sequential approach. This is probably the result of the penalties that were used in the objective function. In the sequential approach the weights used for the objective function of the pre-processing orders are independent of the weights used for the packaging orders. This is not possible in the MIP-based approach and this seems to result in a worse solution.

Due to the relatively long time it takes to solve these instances and because there is no indication that this method will result in large improvements in the solution, no further tests were performed. Although unpractical, It may be possible to improve the solution when the solver is ran for an even longer amount of time.

8.6 Comparison to the current situation at Heemskerk

In this thesis an approach has been presented to find a good production schedule for a production day at Heemskerk. It is interesting to see how this approach compares to the current way in which production orders are scheduled. Due to the format of the available transaction data, it will not be possible to compare the two approaches for all key performance indicators. However, in this section a number of aspects will be used to show how the sequential approach compares to the current manual approach.

8.6.1 Feasibility of the schedule

The current approach at the company uses a fixed machine assignment and production sequence for all orders as a starting point. Because this schedule is fixed and equal for every day of the week, it very often contains infeasibilities at the start. The most common infeasibilities occur because intermediate products are scheduled too late for usage in the final products. This is usually resolved by either postponing packaging, which introduces idle time, or by changing the production sequence. The latter may result in additional changeovers or cleaning.

Another common infeasibility is the overburdening of work centers. Due to promotions or seasonal changes it may be that a certain production line has more work assigned to it than it can process in a single production day. This issue is usually only resolved at the end of the production day when other production lines have completed their assigned orders. These other work centers will then ‘help’ the overburdened work center by doing a part of its work. Obviously this is rather inefficient. The orders that are rescheduled at the end of the production day might have fit better on another production line earlier in the schedule. However, at that point, that opportunity has already been lost.

The approach presented in this thesis provides schedules that are feasible from the start. The relationship between pre-processing and packaging orders is enforced by a fixed constraint that ensures that each intermediate is available when it needs to be packaged. Furthermore the workload on overburdened work centers is distributed over other production lines right at the start of a production day.

8.6.2 Time involved in scheduling of the process

The current manual approach results in a schedule that requires a large amount of on-the-fly adjustments to the schedule. During every shift there is at least one worker that uses all of his time to adjust the production schedule. All these adjustments also result in efficiency losses. An operator may already be halfway through a changeover to another product when his team leader tells him that the schedule of his work center needs to be adjusted.

The approach presented here will not result in the complete removal of all of these adjustments. Adjustments will always be necessary when machines break down, when demand changes or when unprocessed vegetables are unavailable. However, almost all of the other infeasibilities that require changes to the schedule will be taken care of by the optimization approach. This will result in a large reduction in the number of adjustments and will thus increase the efficiency of the process.

8.6.3 Amount of changeovers

From the transaction data that is available at the company, it is unclear whether idle time in the schedule is the result of infeasibilities, machine breakdowns or the unavailability of unprocessed vegetables. It is however possible to extract the amount and type of all changeovers in a schedule. This makes it possible to compare the changeovers in the schedules of both approaches.

The transaction data of the 7th of July 2013 will be used to see how the current approach compares to the designed sequential approach. For each work center the transactions are extracted separately. This will show the exact type and sequence of the orders that were produced on that machine that day. For comparison the different types of changeovers will have the same penalty scores as were used for the sequential approach. Table 8.28 shows a comparison of the total changeover penalty in the schedule of the sequential approach and the actual schedule of the production day.

8.28 Comparison of changeover penalties of the sequential approach and an actual schedule.

Sum of changeover penalties	*Schedule Sequential Approach*	*Actual schedule*
Pre-processing department	424	620
Packaging department	534	786
Total	958	1406

It should be noted that the shown total changeover for the actual schedule is rather high as it also includes changeovers that are the result of rush orders and rework. The two total numbers can therefore not be exactly compared. Nevertheless the table shows that the schedule that is found using the sequential approach is at least in the same ballpark as the schedules that are currently used at the company. It also shows that there is a lot to be gained by using more sophisticated scheduling methods.

9 Conclusion and Recommendations

9.1 Conclusion

In this thesis, the scheduling problem at the large fresh vegetable production company W. Heemskerk was reviewed. The main objective during the research was finding a way to optimize the connection between pre-processing and packaging orders by assigning production resources to the available work. This objective was pursued by creating a mathematical optimization model for the problem and finding a favorable solution for it in a reasonable amount of time. Three research questions were formulated to further investigate this problem.

The first question involves finding a mathematical formulation for the scheduling problem. In section 6 a continuous-time immediate precedence batch-scheduling model was proposed. This model incorporates all process-specific constraints, e.g., the cross-contamination rules and the timely availability of the intermediate products for packaging. The objective function of the presented model minimizes the total tardiness of the orders, the total changeover penalty and the idle time of the orders.

After a mathematical formulation of the problem was found, an attempt was made to solve an instance of the scheduling problem using the commercial solver *CPLEX.* This attempt proved unfruitful due to the large size of the problem instance. To cope with the large size of the problem, a sequential approach to solve the problem was presented in section 7. This sequential approach splits the problem instance into several smaller instances that can be solved using the commercial solver. In this approach, all packaging batches are scheduled first. Next, the schedule that was found for these batches is used as input to schedule the pre-processing batches. Last the solution of the packaging batches is adjusted to find an overall feasible schedule. This approach greatly reduces the number of constraints and variables in each of the instances. This in turn results in better solution times.
The sequential approach was compared to the MIP-based approach using a relatively small problem instance. This comparison shows that the results that are found using the sequential approach are very similar to the results found using the MIP-based approach. However, the sequential approach requires significantly less time to solve the problem.

The strength of the sequential approach was illustrated by solving an actual real-world problem instance at the case company. The results show that, depending on the stopping criteria used, a good solution to the problem can be found in 1 to 2 hours. To further illustrate the possibilities of the approach, the problem was solved using several different combinations of weight factors in the objective function. This showed that increasing or decreasing the weights associated with these objectives directly influences the objectives. This gives the management of the company the possibility to influence the approach.

The approach presented in this thesis was shown to have several advantages over the current manual approach of the case company. In contrast to the current approach, a schedule generated using the sequential approach will be feasible at the start of the production day. This means that a good connection between the pre-processing and the packaging department is ensured by the way in which the model was formulated. Furthermore, the current approach only reschedules excessive workload on production lines at the end of the production day. The method presented in this thesis will evenly distribute the workload in the available production time in the initial schedule. This means that a lot more, potentially better, scheduling options are considered in the process. To conclude, the sequential approach most likely results in a considerable reduction in the amount and the severity of the required changeovers on production lines.

9.2 Recommendations

Several recommendations for further research can be given.

In the current approach, minimizing the sum of all starting times enforces the minimization of the idle time. This is inefficient because a large part of this sum is not influenced by the idle time but by the constraints of the problem. A starting point for improvement could be to maximize the number of orders that have their finishing and starting times in common. This measure would be closer related to the amount of idle time.

Each production day at the case company is a separate problem instance. Because the variety and the sizes of the orders on the same day in different weeks are often very comparable in nature, it may be useful to use the solutions from these days as starting points for the problem instances of future production days. This could further reduce the solution time required to solve a problem instance. It will also be interesting to see how much these schedules differ from each other and which aspects of the problem result in those differences.

It was shown that the stopping criterion used for the solver greatly influences the required solution time. A more intelligent stopping criterion could be formulated by further investigating the results when the solver is stopped after a shorter amount of time. It may well be that when a solution is quickly accepted in the first two steps of the sequential approach, it will not result in a big reduction in the quality of the solution. A start for this was made in section 8.2.4 but further tests would be required to verify the idea that stopping the solver sooner may not result in significantly worse solutions. After all, only little time is required to find a feasible solution in each step of the approach.

References

Akkerman, R., & Donk, D. v. (2009). Analyzing scheduling in the food-processing industry: structure and tasks. *Cognition, Technology & Work vol. 11*, 215-226.

Baumann, P., & Trautmann, N. (2013). A continuous-time MILP model for short-term scheduling of make-and-pack production processes. *Internation journal of Production Research*, 1707-1727.

Christodoulos, F. (2005). Mixed Integer Linear Programming in Process Scheduling: Modeling, Algorithms, and Applications. *Annals of Operations Research 139*, 131-162.

Entrup, M. L., & Günther, H. (2005). Mixed-Integer Linear Programming approaches to shelf-life-integrated planning and scheduling in yoghurt production. *Internation Journal of Production Research 43:23*, 5071-5100.

Floudas, C., & Xiaoxia, L. (2005). Mixed Integer Linear Programming in Process Scheduling: Modeling, Algorithms and Applications. *Annals of Operations Research 139*, 131-162.

Garey, M., & Johnson, D. (1979). A Guide to the Theory of NP-Completeness. *Computers and Intractability*.

Gary, M., Johnson, D., & Sethi, R. (1976). The Complexity of Flowshop and Jobshop Scheduling. *Mathematics of Operations Research, Vol. 1, No. 2*, 117-129.

Gayialis, S., & Tatsiopoulos, I. (2009). Design of an IT-driven decision support system for vehicle routing and scheduling. *European Journal of Operational Research Vol. 152, No. 2*, 382-398.

Grossmann, I., Quesada, I., Raman, R., & Voudouris, V. (1996). Mixed-Integer Optimization Techniques for the Design and Scheduling of Batch Processes. *Batch Processing Systems Engineering*, 451-494.

Herrmann, J. (2006). *Handbook of Production Scheduling.*

Kjellsdotter, L. (2010). User experiences of Advanced Planning and Scheduling (APS) systems in production scheduling.

Kopanos, G., & Puigjaner, L. (2010). Optimal Production Scheduling and Lot-sizing in Dairy Plants: The Yoghurt Production Line. *Industry & Engineering Chemistry Research vol. 49*, 701-718.

Kopanos, G., & Puigjaner, L. (2011). Production Scheduling in Multiproduct Multistage Semicontinuous Food Processes. *Industrial & Engineering Chemistry Research vol. 50*, 6316-6324.

Kopanos, G., & Puigjaner, L. (2011). Resource-constrained production planning in semicontinuous food industries. *Computers and Chemical Engineering vol. 35*, 2929-2944.

Kopanos, G., & Puigjaner, L. (2012). Efficient mathematical frameworks for detailed production scheduling in food processing industries. *Computers and Chemical Engineering vol. 42*, 206-216.

Kopanos, G., Mendez, C., & Puigjaner, L. (2010). MIP-based decomposition strategies for large-scale scheduling problems in multiproduct multistage batch plants: A benchmark scheduling problem of the pharmaceutical industry. *European Journal of Operational Research*, 644-655.

Kopanos, G., Puigjaner, L., & Geogiadis, M. (2012b). Efficient Mathematical frameworks for detailed production scheduling in food processing industries. *Computers and Chemical Engineering 42*, 206-216.

Lenstra, J., & Rinnooy Kan, A. (1977). Complexity of machine scheduling problems. *Annals of Discrete Mathematics 1* , 343-362.

Lin, C.-H., Hwang, S.-L., & Min-Yang Wang, E. (2007). A reappraisal on advanced planning and scheduling systems. *Industrial Management & Data Systems 107*, 1212-1226.

Marinelli, F., & Nenni, M. (2007). Capacitated lot sizing and scheduling with parallel machines and shared buffers: A case study in a packaging company. *Annals of Operation Research vol. 150*, 177-192.

Mendez, C., & Cerdá, J. (2002). An MILP-based approach to the short-term scheduling of make-and-pack continuous production plants. *OR spectrum vol. 24*, 403-429.

Mendez, C., & Cerda, J. (2006). State-of-the-art review of optimization methods for short-term scheduling of batch processes. *Computers and Chemical Engineering 30*, 913-946.

Mendez, C., Cerda, J., & Grossmann, I. (2006). State-of-the-art review of optimization methods for short-term scheduling of batch processes. *Computers and Chemical Engineering 30*, 913-946.

Pekny, J., & Reklaitis, G. (1998). Towards the Convergence of Theory and Practice: A Technology Guide for Scheduling/Planning Methodology. *Proceedings of the Third International Conference on Foundations of Computer-Aided Process Operations*, 275-309.

Reklaitis, G. (1992). Overview of Scheduling and Planning of Batch Process Operations. *Batch Proces Systems Engineering*.

Ruiz, R., & Stutzle, T. (2008). An iterated greedy heuristic for the sequence dependent setup times flowshop problem with makespan and weighted tardiness objectives. *European Journal of Operational Research 187*, 1143-1159.

Shah, N. (1998). Single- And Multisite Planning and Scheduling: Current Status and Future Challenges. *Proceedings of the Third International Conference on Foundations of Computer-Aided Process Operations*, 75-90.

Stadtler, H., & Kilger, C. (2005). *Supply Chain mangement and Advanced Planning-Concepts, Models, Software and Case Studies.* Berlin: Springer.

Stefansson, H., & Sigmarsdottir, S. (2011). Discrete and continuous time representations and mathematical models for large production scheduling problem: A case study from the pharmaceutical industry. *European Journal of Operational Research 215*, 383-392.

Wauters, T., & Verbeeck, K. (2012). Real-world production scheduling for the food industry. *Engineering Applications of Artificial Intelligence*, 222-228.

Appendices

A. List of Figures

B. List of Tables