MASTER THESIS

Pioneering with Mixed Integer Linear Programming as a Solution to a Two-Stage Hybrid Flow Shop with No Buffer Capacity

A case study for confectionary manufacturer Brynild

Author
C.M.S. van der Valk (Charlotte)

Examination committee

Gréanne Maan-Leeftink, PhD	*University of Twente*
Marco Schutten, PhD	*University of Twente*
Hans de Man, MSc	*SINTEF*
Mathias Holm, MSc	*Brynild Gruppen AS*

Education Institution
University of Twente
Faculty of Behavioural Management and Social sciences
Section of Industrial Engineering and Business Information Systems

Educational Program
MSc. Industrial Engineering and Management
Specialization: Production and Logistics Management
Orientation: Supply Chain and Transportation Management

April 2021

Brynild
Smaksgleder siden 1895

SINTEF

UNIVERSITY OF TWENTE.

MANAGEMENT SUMMARY

Brynild is a Norwegian manufacturer of confectionary. We perceive Brynild's current production scheduling processes as inefficient; Brynild creates their current schedules manually using experience-based techniques, and lacking insight in whether the resulting schedules deliver good or poor performance. Our data analyses of the current scheduling method indicates that while the bottleneck, the drying area, is almost always occupied, namely 89% to 97% of the time, it is far from fully utilized, solely 40% to 57% of the time. Therefore, the research goal of this thesis is:

'*Develop a scheduling method that improves the scheduling of the production orders under consideration of multiple process constraints*'

System settings: two-stage hybrid flow shop with zero buffer capacity

We classify the production line under consideration as a two-stage hybrid flow shop. This hybrid flow shop consists of the following stages:

- Stage 1: a single line including sequential processes such as, set-up of the line, cooking, and moulding of the confectionary;
- Stage 2: 5 parallel drying cabinets. Stage 2 starts when the moulding of the confectionary starts and ends when the drying cabinet finishes.

Stage 1 and Stage 2 overlap during the moulding time since both the single line and one of the parallel machines need to be available during moulding.

The two-stage hybrid flow shop has the following characteristics: 1) non-identical parallel machines (drying cabinets) in terms of processing time and capabilities of drying certain products; 2) sequence dependent set-up times; 3) production time windows; and 4) zero buffer capacity between the 2 stages. The fourth characteristic, scheduling with zero buffer capacity, is only briefly analysed in the existing literature, and with this research we contribute to the literature on this relatively new problem configuration.

Methods: MILP with three MILP-based heuristics

We base our MILP model on input data that we structure in such a way that jobs consist of 1 or more intermediates that can dry together and, in terms of quantity do fit in 1 drying cabinet. We obtain this input data by pre-processing the data, using a heuristic that joins 2 intermediates in 1 job when this fits in the drying cabinet. And in the case that the intermediates' quantity exceeds the capacity of the drying cabinets, the heuristic divides the quantity of the intermediate evenly over the minimum number of jobs.

For the two-stage hybrid flow shop problem, we develop a mixed integer linear programming (MILP) scheduling model that minimizes the make span. We also develop a second objective function for this model that minimizes the maximum end time for the first stage, since only this stage involves human activities. We conduct multiple experiments for which we use 7 of Brynild's production weeks.

For certain weeks, the MILP requires more than 1 night (16 hours) to compute; this is too long for practical purposes. Therefore, we develop 3 heuristics with shorter computation time that we base on our MILP:

- Our MILP without sequence dependent set-up times;
- Our MILP that we decompose in an assignment and a sequencing MILP, which we solve sequentially;
- Our MILP with a maximum computation time of 10 minutes.

The results: significant improvements in maximum make span

In comparison to Brynild's original week schedules, our schedules perform better:

- When minimizing the end-drying time, the last drying cabinet finishes significantly earlier ($p<0.05$); *on average 38 hours, with peaks up to 49 hours earlier*;
- When minimizing the end Stage 1 time, the last job on Stage 1 finishes significantly earlier ($p<0.05$); *on average 17.5 hours, with peaks up to 39 hours earlier*.

Brynild works with shift schedules, which inform managers how much personnel to hire per day of the week. Our experiments show that:

- We schedule *75% of the originally 3-shift schedules in 29% less time*;
- We schedule *33% of the originally 2-shift + Saturday schedules in 10% less time*.

The heuristic with a maximum computation time of 10 minutes, performs the best for both objectives:

- When minimizing the end-drying time, the heuristic *re-creates 86% of the optimal schedules*;
- When minimizing the end Stage 1 time, the heuristic performs *the most stable* in terms of deviation from the optimal solution, with a variance of 0.826.

To show that our methods work in practice, we schedule weeks 7 and 8 of 2021 with live data from Brynild, to analyse whether the schedules we create can be adopted in practice. After several iterations with the scheduler's insights, the operators of the production line gave feedback on the practical feasibility and responded positively that they '*find it hard to believe that the scheduling is done by a model*´.

Besides Brynild's case study we evaluate the general capabilities of our main MILP using the end-drying objective. We observe an exponential increase in computation time for all 3 types of instances (short, long, or mixed drying times) when the number of jobs increase. When the number of parallel machines increases, the number of jobs that we can schedule within 16 hours, decreases. Next, we examine 8 two-stage hybrid flow shops varying in including production time windows, sequence dependent set-up times, and non-identical parallel machines, that share the characteristic of zero buffer capacity between the 2 stages. We conclude that models without sequence dependent set-up times compute much faster than models with sequence dependent set-up times. Next, in almost all cases the models without production time windows compute faster than the models with production time windows. Finally, including identical parallel machines results in both increase and decrease in computation time. After further evaluation, we hypothesise that when the process time becomes the same for the parallel machines, the model has a faster computation time; and when all parallel machines can process all jobs, the solution space becomes larger and the computation time longer.

Recommendations

We recommend Brynild to use the Stage 1 objective 3 weeks in advance to create a week schedule and to determine the minimum number of shifts. Brynild should consider this minimum number of shifts while selecting the shift schedule. Thereafter, we recommend Brynild to use the MILP with the end-drying time objective, in which they incorporate the selected shift schedule as a constraint. If Brynild wants to do scenario testing, to analyse whether the number of intermediates or the quantity of the intermediates can increase within the week, we recommend using the 10-min run heuristic.

In our analysis, we make some assumptions that offer promising possibilities for future research. For example, including in the MILP, batch composition, as this characteristic opens many possibilities for scheduling results and enhances the literature. Furthermore, our experiments show that the model computation time highly depends on the number of parallel machines and jobs. Therefore, when Brynild decides to expand its production line with another drying cabinet or by increasing the number of jobs, we expect that not all weeks can be scheduled within 16 hours when using the developed MILP. In this case, we recommend Brynild to use the MILP and interrupt the MILP after 16 hours, as the model has most likely found the (near) optimal solution within this time frame. This thesis offers Brynild and other companies with likewise problems, a method to greatly improve their scheduling method.

PREFACE

In September 2020, I started my journey in Norway as a guest researcher at SINTEF. I was welcomed with open arms to perform research for my graduation assignment of the master's degree Industrial Engineering & Management at the University of Twente. SINTEF supervised me during this period while I performed research for Brynild Gruppen AS. I visited Brynild's factory for one week during my stay in Norway. During this week, Brynild was a great host, and I was given the freedom to ask questions to anyone within the factory. Beside this week, I could always ask my questions by mail or during online meetings.

I would like to take the opportunity to express my gratitude to the people who helped me realize this thesis.

First of all, I would like to thank Gréanne Leeftink and Marco Schutten from University Twente for their supervision. Based on their constructive feedback and support, I was challenged to broaden my view on the subject and their ideas gave me new insights for my thesis. They also helped me with formulating the thesis, which makes the thesis as it is now. I would like to thank Sonja Borst as well as she was my first contact person from the University Twente for this graduation assignment.

Furthermore, I would like to thank Hans Torvatn for facilitating a workspace at SINTEF. My special thanks go to Hans de Man, who made it possible for me to move to Norway and have an amazing experience during my stay in Trondheim. Hans de Man was my daily supervisor at SINTEF and he was always ready to help. Next, I like to thank Mathias Holm for the facilitation of my visitation week at Brynild, and opening Brynild's doors for me. Thereafter, thank you Eirik Blå for answering my questions and for your help to make the model represent practice as much as possible. I also like to thank the people in Brynild's operation for the given tours through the factory and explaining the processes.

Last, I like to thank my family and friends for their encouragement and support during this period. I look back with great pleasure at my time in Norway and working on this thesis.

Charlotte van der Valk
Soest, April 2021

LIST OF ABBREVIATIONS

Abbreviation	**Explanation**
AP	Assignment Problem
BH	Backward Heuristic
CODP	Customer Order Decoupling Point
ERP	Enterprise Resource Planning
FH	Forward Heuristic
FJSP	Flexible Job Shop Problem
GA	Genetic Algorithm
GAP	Generalized Assignment Problem
GQAP	Generalized Quadratic Assignment Problem
HFS	Hybrid Flow Shop
IA	Immune Algorithm
IABC	Improved artificial bee colony
IG	Iterative Greedy
ILS	Iterative local search
JSP	Job Shop Problem
KPI	Key Performance Indicator
LSAP	Linear Sum Assignment Problem
MILP	Mixed Integer Linear Programming
MPS	Master Production Schedule
NEH	Nawaz-Enscore-Ham
PVNS	Parallel Variable Neighbourhood Search
QMKP	Quadratic Multiple Knapsack Problem
RKGA	Random Key Genetic Algorithm
RPD	Relative Percentage Deviation
SA	Simulated Annealing
SDST	Sequence dependent set-up time
SKU	Stock Keeping Unit
TS	Tabu Search
VNS	Variable Neighbourhood Search
WIP	Work in progress

TABLE OF CONTENTS

1 INTRODUCTION

The purpose of this master thesis is to explore multiple scheduling methods for a sequence dependent production line. With a literature study and some new concepts, we want to provide Brynild Gruppen AS advice on the scheduling of their production line. Section 1.1 describes the organizational context. Section 1.2 introduces the problems experienced by Brynild, and in Section 1.3 we present the research goal and articulate our approach to solve the problem. Finally, Section 1.4 presents the research scope.

1.1 Organizational context of SINTEF and Brynild Gruppen AS

This research is performed in collaboration with SINTEF and commissioned by the company Brynild Gruppen AS. SINTEF has been given the task to optimize Brynild Gruppen's production line. We introduce both companies further in Section 1.1.1, Section 1.1.2, and Section 1.1.3.

1.1.1 SINTEF

SINTEF, founded in 1950, is one of the largest independent research organisations within Europe. The headquarter of SINTEF is located in Trondheim, Norway. There are 2000 employees with 75 different nationalities active at SINTEF. They deliver applied research, technology development, knowledge, innovation, and solutions for 3600 large and small customers from around the world. This makes that SINTEF excels in over 400 research areas, varying from ocean space to outer space and everything in between. SINTEF's vision is "Technology for a better society" (SINTEF, 2020). We conduct this specific research at the department of technology management within the group Learning & Decision support in Trondheim.

1.1.2 Brynild Gruppen AS

Brynild Gruppen AS is one of Norway's largest confectionary manufacturers and family-owned by the 4^{th} and 5^{th} generation. They have roots going back to 1895. Brynild gruppen Holding AS is the parent company of various brands and operates in fast-moving consumer goods, such as chocolate, confectionary (sweets) and snacks (nuts and dried fruits).

Den Lille Nøttefabrikken, Brynild, Minde Sjokolade, Dent, and St. Michael are Brynild Gruppen's largest brands (see Figure 1.1). The market position of Brynild Gruppen ranges from being a minor challenger (in chocolate) to being a market leader (in nuts). Next to this, Brynild Gruppen is also a distributor for the German company Beiersdorf, which is best known for their brands NIVEA, Labello and Hansaplast.

Brynild Gruppen is responsible for all the value adding processes; starting from the product development, purchasing, the logistics and production up to the sales and marketing of their own brands, that consist of approximately 350 Stock Keeping Units (SKUs). They have 200 employees working for them, and the annual turnover is 750 million NOK, i.e., 70.6 million euros. 90% of the sales are originated from Norway and the other 10% from other Nordic countries (Brynild Gruppen, 2020).

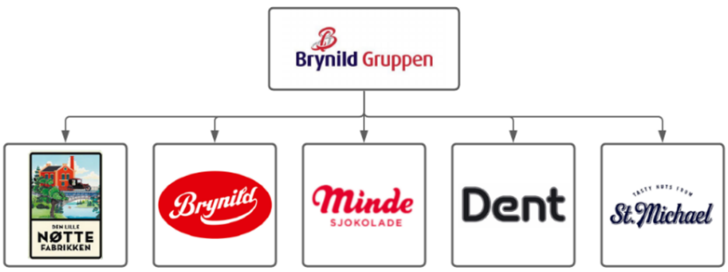

Figure 1.1: Brands Brynild Gruppen

1.1.3 Production line Brynild Gruppen AS

Brynild Gruppen AS has 1 production plant in Fredrikstad, Norway. This production plant has 3 production lines that produce all the different products of the various brands. We focus on the confectionary production line, which produces the Brynild and Dent brands. This production line produces all the various confectionary products, as we show in Figure 1.2.

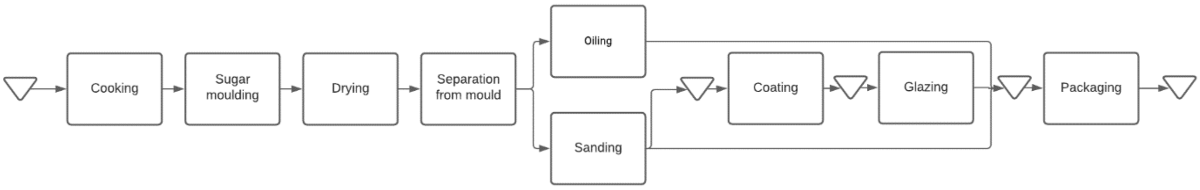

Figure 1.2: Production line confectionery

Figure 1.2 shows the processes in the squares, and the work in progress (WIP) in the triangles. The process starts with the input of raw materials. With these raw materials, Brynild Gruppen prepares the candy, which is called 'intermediate' until it is packaged, by cooking the raw materials. After cooking the mass, Brynild Gruppen adds the flavour and colour. When Brynild Gruppen finishes this process step, they mould the cooked substance in the right shapes. After the moulding process, the intermediates need to dry. The drying process requires different humidity for different types of confectionary, such as temperature and duration. After drying, the moulding machine separates the confectionary from its mould, and two process options arise. Option 1: oiling; Option 2: sour or sugar sanding. One of these two options must always occur. The oiled intermediates and most of the sanded intermediates go straight into the packaging station. Only some of the sanded intermediates require coating and glazing before Brynild can package them. Within this production line buffer points occur, namely before cooking (the raw material), and after sanding & oiling, coating, glazing, and packaging.

1.2 PROBLEM STATEMENT

Brynild Gruppen AS, to we refer to as Brynild from now on, has several strong competitors such as Mars, Chupa Chups, Kellogg's, Haribo, Lindt, Tic Tac, Nestlé, Ferrero, Kinder, Milka, Orkla (Nidar), Mondelez, and more. Due to these strong competitors, Brynild focuses on product innovation and building brands. To stay ahead of the competitors, Brynild aims to continuously optimize the following operational targets:

- Achieve high service levels;
- Minimize inventory levels;
- Maximize resource utilization;
- Minimize obsolescence and its risks;
- Simplicity in production planning.

The demand at Brynild has grown with 25% over the last 2 years. A consequence of the growing demand is that Brynild is not able to produce enough products. Norwegian grocery retail demands an average of 98% in full and on time from food suppliers. It is however difficult for Brynild to achieve these high service levels and therefore Brynild has to say 'no' more often. This results in dissatisfaction with customers (retail stores), which in turn causes reputation damage for Brynild.

1.2.1 Problem cluster

To investigate how we can improve the production capacity of Brynild, we create a problem cluster in order to identify the cause and effect relationships that lead to the core problem(s) (Heerkens & Winden, 2012). As Brynild focuses on different operational targets we mention the ones that occur in the problem cluster:

- High service levels are difficult to achieve as meeting the order due date becomes more difficult;
- Bottleneck (drying area) resources are not fully utilized;
- The process of creating the schedules is not generic, therefore there is no simplicity in the production planning and scheduling.

We explain the problem cluster in this section and Figure 1.3 presents an overview of the problem cluster. From Figure 1.3 we observe 3 final problems: reputation damage, the scheduling requires a lot of time and high risk in relation to absence for sickness and/or holidays.

The high risk that comes with sickness and holidays, is a consequence of having only 1 employee who exactly knows how to schedule the confectionary production line. This makes the process of scheduling production orders non-generic, as the scheduling is done experience-based. Therefore, it is difficult for other employees to repeat this process, which makes Brynild reliable on 1 specific employee. Another problem that arises from working experience-based is that obtaining the schedules takes a lot of time. The process is often trial and error, which is very time consuming, without knowing if the schedule is as efficient as it could be. The reason that Brynild conducts the sequence and assignment scheduling manually and experience-based, is because there is no advanced scheduling method within Brynild's Enterprise Resource Planning (ERP) system.

Brynild indicates reputation damage as the most serious problem that arises. Dissatisfied (future) customers are the cause for the reputation damage. Currently the reputation damage is minor as Brynild 'only' has to reject new clients and promotions. However, this is a temporary escape to deal with the capacity constraint Brynild experiences. Another reason for the dissatisfaction of the customers is that not all orders meet the order due date. Not being able to meet the due dates is another consequence of the higher demand; as demand is higher than the production capacity, Brynild cannot produce enough products on time. The limitations of the production capacity have two main causes:

1. There is limited amount of floor space for the WIP;
2. The drying area is often the bottleneck, which stagnates the whole production line.

There is not enough free space on the Brynild's work floor, which causes the limited amount of floor space for the WIP. Not enough free space indicates that the production location is not big enough for how Brynild currently produces the products.

One of the reasons that the drying area is the bottleneck is, according to Brynild, the availability of drying cabinets during a week that produces 24 hours a day. That is why Brynild is in the process of buying and installing an additional drying cabinet. However, even with the additional drying cabinet Brynild indicates that the drying area is likely to remain the bottleneck.

The other reason that the drying area is the bottleneck, is because Brynild does not fully utilize the drying cabinets. Often, the drying cabinets are full, but not fully used. Full, but not fully used, refers to two different aspects. On the one hand, the cabinet is drying, however it is only half full due to batch sizes that do not consider the entire capacity of the drying cabinet. On the other hand, the products stay longer in the drying cabinet than necessary. Brynild only empties the drying cabinet when a new batch enters the cabinet. These two reasons make the effective utilization of the drying area lower than it can be. This inefficient use of the drying cabinets is due to experience-based assignment and sequencing. The reason for applying experience-based assignment and sequencing is, as mentioned earlier, that there is no advanced scheduling method in place.

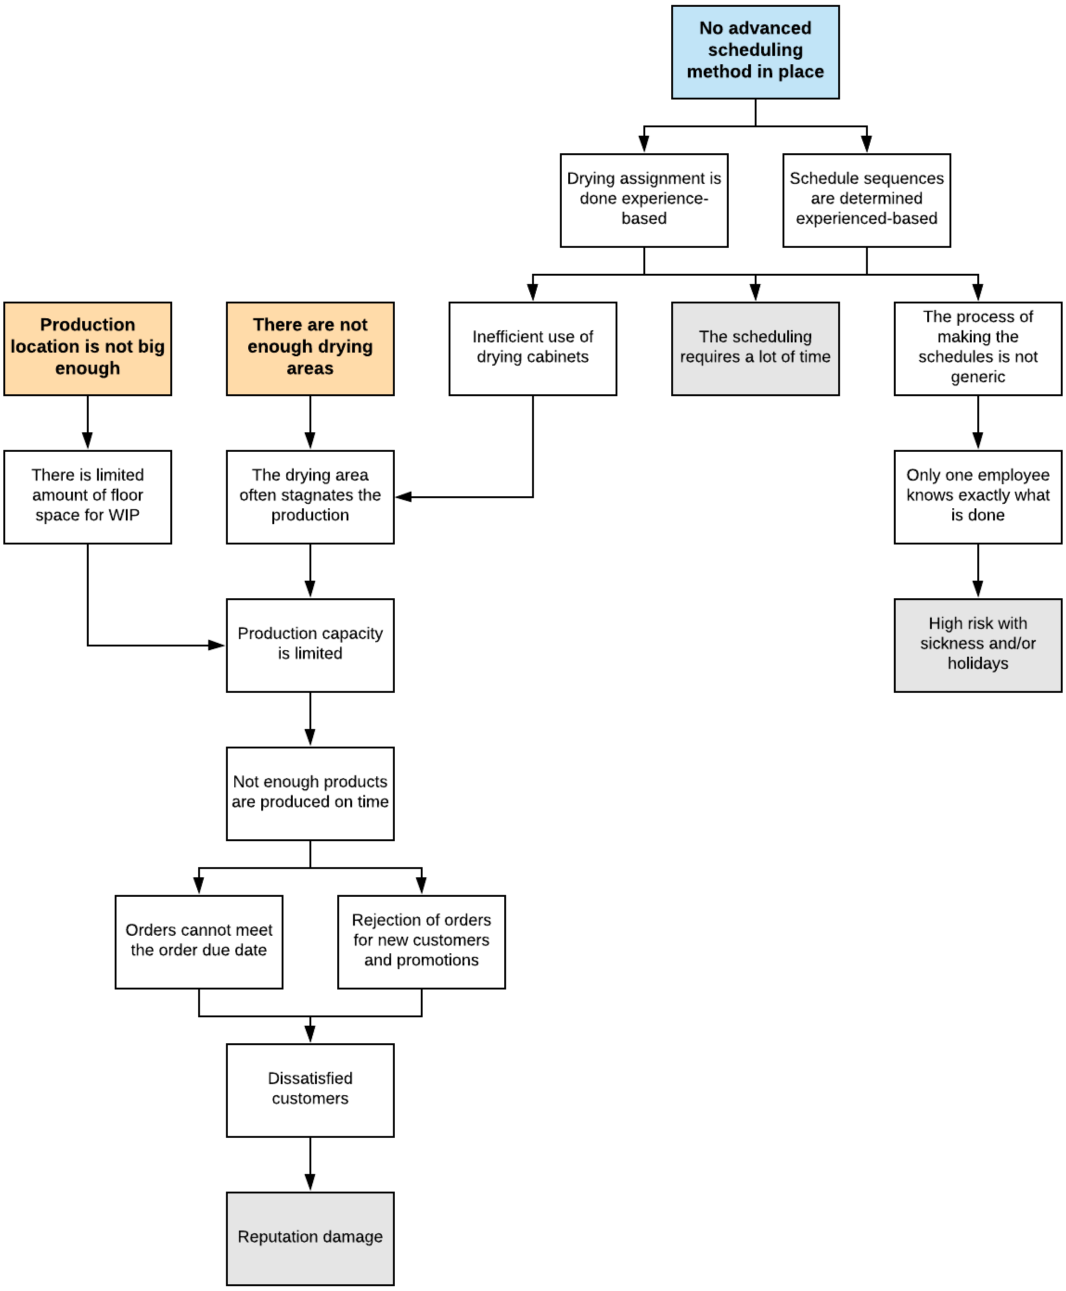

Figure 1.3: Problem cluster

1.2.2 Core problem

We find the core problem in the problem cluster by following the rules of thumb (Heerkens & Winden, 2012). The problem cluster shows 3 problems that do not have a further cause themselves and that we could influence in the context of this research, namely:

- "Production location is not big enough";
- "There are not enough drying areas";
- "No advanced scheduling method in place".

There is one problem that we can solve without altering Brynild's production space, and, which has the biggest influence on the rest of the cluster "no advanced scheduling method in place". When Brynild conducts the scheduling using a better method, the probability is high that the production line obtains more production capacity. Next to more production capacity, an advanced scheduling method comes with a more generic approach, which makes it possible for the production schedulers of different lines to be interchangeable. On top of that, the scheduling requires less personnel worktime, as the method does most of the work. To make this problem more quantifiable, we focus on the following Key Performance Indicators (KPIs):

- The percentage of production orders that meet the planned due date;
- The occupation of the bottleneck; the drying area;
- The utilization of the bottleneck;
- The utilization of the bottleneck in comparison with the occupation;
- The make span of the production week;
- The changeovers set-up times.

The KPIs above give a clear grasp on the magnitude of Brynild's scheduling problem. Next to getting a better insight, we also use (most of) these KPIs to compare the performance of the schedules we create in this research.

1.3 RESEARCH GOAL

This section presents the main research goal and the corresponding research questions.

1.3.1 Main research goal

Together with Brynild Gruppen AS, we desire to improve the scheduling of the production line to free up valuable production capacity, which improves the throughput of the production line. This generic method should require a minimum amount of knowledge and time of the scheduler.

The research goal is:

> *'Develop a scheduling method that improves the scheduling of the production orders under consideration of multiple process constraints'*

This research goal results in the following main research question:

> *'How should we construct the scheduling method for Brynild Gruppen AS, such that the production line can realize a higher throughput?'*

1.3.2 Research questions

To answer the central research question, we formulate the following sub-research questions. Each chapter of this thesis answers one of the questions.

1. **What does the current situation at Brynild look like?**
   In Chapter 2 we give an overview of characteristics and conditions of the production line and the scheduling process. This overview gives a better understanding of the underlying problems we describe in Section 1.2. To obtain the information we need, we conduct a series of interviews with the supply chain manager, the scheduler, and the production manager. Next to interviews, we do observations at the production area. We collect associated documentation, such as schematic overviews and historical data of the production and scheduling process. Then, we analyse the performance of the current situation on the KPIs we mention in Section 1.2.2. With this information we evaluate the current situation, which provides us the benchmark for the improvement methods.

2. **Which methods are described in literature regarding the scheduling of a production line similar to Brynild's?**
   To answer the second research question, we perform a review and analyse the relevant scientific literature in Chapter 3. We focus on scheduling methods that tackle the assignment problem and the sequencing problem. We describe the similarities, differences, and applications of the scheduling problems that are similar to Brynild's problem. Next to that, we review multiple opportunities on how to solve the problems. Finally, we discuss, which methods are applicable for Brynild and, which method is the most suitable for our research.

3. **How can we develop a scheduling model that improves the throughput of the production line?**
   Chapter 4 presents a mathematical formulation of the problem we describe in Chapter 2. First, we decide, which variables and parameters we consider while formulating the scheduling problem. The scheduling problem beholds the characteristics and conditions of the production process. We evaluate multiple objectives that could lead to the focus of the scheduling method; obtaining a higher throughput.

4. **How does the proposed scheduling method perform?**
   In Chapter 5, we test the performance of the model against practical instances. We also study the differences between our method and the current way of scheduling. This comparison is in the form of a case study that we base on historical data. After creating the new schedules, experts analyse our method to verify them based on the practical use for Brynild. Thereafter, we examine the performance of the general use of the model we develop.

Chapter 6 contains the conclusions and recommendations of this research.

1.4 SCOPE OF THE RESEARCH

We base our scheduling on the confectionery production line in Fredrikstad. This implies that we exclude the other production lines, for nuts and chocolate, of Brynild Gruppen AS from this research. The planning needs to provide the intermediates and the weight of the intermediates as production orders for each specific week. This planning considers the variability from outside the production line; like the availability of raw materials, maintenance etc..

2 CURRENT SITUATION

In this chapter we answer the first research question: '*What does the current situation at Brynild look like?*'. Therefore, we give a detailed overview of the confectionary production process in Section 2.1. In Section 2.2 we describe the scheduling process including all challenges it encounters. In Section 2.3 we discuss the current performance of the confectionary line, and in Section 2.4 we give a conclusion of this chapter.

2.1 BRYNILD'S CONFECTIONARY PRODUCTION PROCESS

Brynild produces 38 different intermediates on their confectionary production line. These intermediates are divided into 6 product families. Brynild bases these families on the drying characteristics. From these 38 intermediates, 7 have to go past the Drage Sukker line, the other 31 intermediates go directly to packaging, see Figure 2.1.

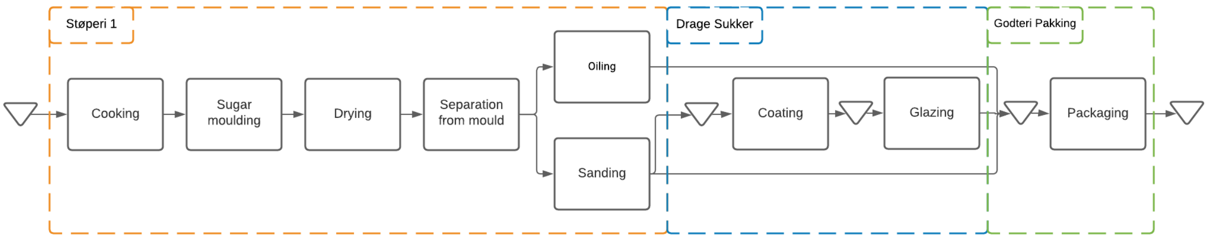

Figure 2.1: Processing stages sugar confectionary line

2.1.1 Cooking and Moulding

The first two processing stages are cooking and moulding. In the first stage, Brynild mixes and cooks the raw materials. After cooking, the mass transfers through pipes to the moulding machine. Cooking starts 30 minutes before moulding. Within each batch the colouring and taste may vary as this has no further consequences for the next processes. At the moulding machine, Brynild sprays the sugar confectionary into wooden trays that contain corn flour pre-formed beds in the shape of the correct intermediate. The trays can contain between 96 and 800 pieces, depending on the product type. After moulding, Brynild stacks the trays on pallets and automatically transfers the pallets to the drying area. One pallet can contain 150 wooden trays, independent of the product type, see Figure 2.2. On average, Brynild can mould 27 trays every minute.

Figure 2.2: Transportation to the drying area

2.1.2 Drying and Separation from mould

Directly after cooking and moulding, the drying process starts. The process, starting from cooking until the intermediates enter the drying cabinets is a continuous process, which does not have any buffer points in between. The drying processing stage has 7 cabinets with different characteristics and capacities. There are 3 various drying cabinets: Catelli, Dynaflow and Lateskap (passive cabinet). The Catelli and Dynaflow are modern drying cabinets, which have shorter throughput time than the Lateskap, the older cabinet. The drying temperature ranges from room temperature 20°C to 65°C. The drying time can vary between 24 hours to 94 hours, depending on the drying temperature and on the type of intermediate. As the humidity in summer, especially in August, is higher than in winter, some products have different winter and summer drying times. Since Brynild is in the process of buying dehumidifiers, we do not take this variation in drying times into account in our research.

There is 1 product family, in Norwegian called 'Familie D', that can only use the new drying cabinets to preserve the quality of the products. We present the capacity per drying cabinet in Figure 2.3.

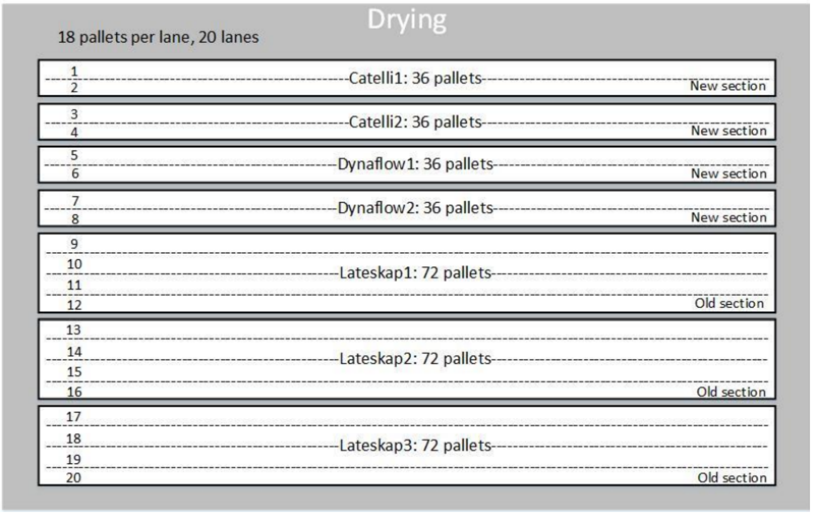

Figure 2.3: Capacity drying cabinets

All drying cabinets together have a capacity of 20 lanes. Each lane can store 18 pallets. Therefore, in theory, a total of 54,000 trays could dry at the same time. As Brynild states the quantity of each intermediate in kilogram, the data for the capacity of the drying area for each intermediate is also in kg per lane, see Appendix A.

Some intermediates from the same family can dry at the same time in the same drying cabinet. In this case, the drying cabinet has to be available when the first intermediate starts moulding. The drying can only start when the final tray from the second intermediate enters the drying section, as the whole drying cabinet locks and heats simultaneously. Not all products from the same family can dry together, see Table 2.1.

Table 2.1: Product families

Product family	# of intermediates	Same drying temp and time
Familie A	7	1,1,5
Familie B	1	1
Familie C	1	1
Familie D	10	2,3,5
Familie K *(Kaldstøp)*	16	1,5,10
Familie	3	1,2

We write all the product family names in Table 2.1 using their Norwegian spelling. When intermediates have the same drying temperature and the same drying time, they can dry in the same cabinet. We summarize this in the last column, for example in Familie A, 5 intermediates could dry together and the other 2 need to dry individually.

The drying cabinets have different characteristics. The Catelli cabinet is heated, circulated, and ventilated. The Dynaflow cabinet is only heated and circulated, and the Lateskap cabinet is only heated. Achieving the right temperature and cooling takes more time in the Lateskap than in the new sections, which lead to longer total drying time. We present the impact of the new drying cabinets in an example for Jordbærfisker (strawberry fish) in Figure 2.4, which shows that this intermediate needs 92 hours drying in the old drying cabinet, and 56 hours in the new drying cabinet.

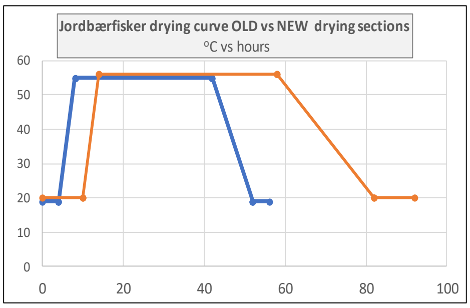

Figure 2.4:The old vs new drying cabinets

After the intermediates finish drying, they are sent back to the moulding machine. At the moulding machine, the intermediates separate from their trays, we call this demoulding. The trays are already in the moulding machine where Brynild uses the trays immediately for a new batch of intermediates. There is a limited number of trays, namely the 54,000 that fit in the drying areas. There is no additional storage for the trays. Therefore, the trays that leave the drying cabinet return to the exact same drying cabinet. As moulding and demoulding is a continuous process it takes approximately the same amount of time, which is 27 trays per minute on average. After the separation from their moulds, the intermediates transfer to the next stage: oiling or sanding.

2.1.3 Oiling or Sanding

In this stage the intermediates convey to either the oiling drum or the sanding drum. There are 6 different sanding/sugaring types of which 3 are sugar free. To make the sanding/sugar stick to the intermediates, steam is shortly applied to the surface just before the sanding drum. Brynild uses only one type of oil/wax is to give the intermediate a shiny surface. Directly after oiling or sanding, Brynild fills plastic boxes of approximately 8 to 10 kg each with the intermediates. This process requires about 30 minutes. Next, the intermediates have to rest in the boxes for 24 hours before further processing can start. 18% of the intermediates that come from the oiling drum need further processing and go to the coating and glazing stage, all the other intermediates go straight to packaging.

2.1.4 Coating and Glazing

In this stage the intermediates go through the coating and glazing processes. This is the process where additional flavour and/or colour is added in parallel rotating drums. Approximately 40% of the intermediates' final weight stems from the coating process. Brynild has 16 coating drums of which 8 are pear shaped, the others are apple shaped. The pear-shaped drums have 5% more capacity per batch than the apple shaped ones. The intermediates of Familie D are only fit for the pear-shaped drums. The capacity of coating is 1400 kg/shift for sugar containing intermediates, or 700 kg/shift for sugar free intermediates. This capacity doubles when 2 operators are working simultaneously. The coated products need to rest for 24 hours before they can be glazed.

All intermediates that go through the coating process must go through glazing process as well. With glazing, some small amount of wax is added to polish the surface and make the intermediates ready for packaging. The glazing is done in one of the 3 bigger drums that. These drums have a capacity of 3500 kg/shift. All glazed products need ripening, i.e., time to mature. The following rules apply:

- Sugar containing products can be packed 48 hours after glazing;
- Sugar free products can be packed 72 hours after glazing.

The feeding into the drums and emptying the drums into plastic boxes is done manually for both coating and glazing.

2.1.5 Packaging

Brynild packs the intermediates first into consumer packages (F-pak) and then into distribution packages (D-pak). The distribution packages are placed onto pallets and become inventory at the warehouse of Leman, 40 minutes away from Fredrikstad. Three to four times a day a truck leaves to restock the warehouse. There are 4 packaging lines that package different products: Dent, Bulk, Bosch and HGD. Finished SKUs consist of one or several intermediates gathered in one consumer package. Some intermediates are sold both as individual products and as a part of consumer packages, consisting of a mixed set of intermediates. The mixed packages can only start packaging when all intermediates are available.

2.1.6 Work in Progress

Figure 2.1 shows that there are several WIP inventory points. Brynild stores their intermediates in plastic boxes at these WIP inventory points. One pallet can hold 40 boxes, 4 boxes per layer and 10 on top of each other. The boxes contain 8 to 10 kg of intermediates. There are around 8,000 boxes available for the confectionery production line. The floor space in this area is limited. Brynild does not know the capacity of the floor space per buffer point, as all the buffer points use the same floor space.

2.2 PRODUCTION SCHEDULING AT BRYNILD

In this section we describe how Brynild determines their Master Production Schedule (MPS) and how they use the MPS output in the scheduling process.

Brynild uses SAP's ERP system. This ERP system reports the weekly demand for each intermediate. The ERP system considers the lead time and historical demand in the previous year. The output of the ERP system is unevenly distributed over the weeks. To level the demand, the production planner manually changes this weekly demand using a "drag and drop" approach. The planner typically drags demand to an earlier week in the year according to the following priority rules:

1. Seasonal products;
2. New products; per year, there are only 3 time windows for introduction;
3. Packed bags;
4. Bulk products for pick and mix.

These priority rules are based on "not wanting to lose the sale by producing it too late". The production planner considers known factors such as planned maintenance or lack of raw materials. This is how Brynild determines their MPS. The MPS states the production orders that consists of which intermediates, in what week, and the quantity to produce. Brynild uses the information from the MPS in their scheduling process.

The ERP system does not determine the sequence or drying assignment of the production orders. Therefore, the scheduler creates the sequence and assignment schedules manually. The scheduler tries to obtain satisfactory schedules based on experience-based techniques, where he makes use of Microsoft Excel. After the schedule is created, the scheduler prints the schedule, as there is no digital version of the production schedule on the production floor. This print includes the recipe for the intermediates and the bill of material for each SKU. The planner/scheduler makes the planning every week and uses an iterative approach for the upcoming week to make sure that the planning fits in the schedule.

The experience-based techniques for scheduling mainly focuses on the drying cabinets. The two rules of thumb the scheduler considers while scheduling are:

1. The production orders need to fit in the drying cabinets;
2. Try to use the new drying cabinets as much as possible.

Figure 2.5 presents an example schedule with the different drying cabinets on the left, in the upper part the products, and the amount in kg that goes into each of the specific drying cabinets. It also states d, k or n, which stands for d= dag (day), k= kveld (evening) and n= natt (night). The lower part of the schedule shows what to unload from the drying cabinet. Note that most emptied drying cabinets are also the ones that are filled, as this is a continuous process.

Prosessoversikt uke 40	Mandag					Tirsdag					Onsdag				
Støperi 1	Art.nr	Artikkel	Kg/stk		Ordre	Art.nr	Artikkel	Kg/stk		Ordre	Art.nr	Artikkel	Kg/stk		Ordre
Skap 1						106497	Sure Skumfrosker	5300	d	227805	112873	Sure Colaflasker	6300	k	227807
Skap 2						113541	Gompegele	4000	k	227806		Kaffebrun farge kommer mandag			
Skap 3							OBS! Sjekk utbytte til gompemix				112815	Knatter	3600	d	227808
		Pausekjøring dag & kveld pga mengde opptak													
Catelli	111132	Sure Buttons	5000	d	227803										
Dynaflow	111132	Sure Buttons	3500	d/k	227803										
	113496	Sure Tær	3800	k	227804										
Opptak															
Skap 1						113729	Supermixgele	7800	d		106497	Sure Skumfrosker	5300	k	
Skap 2						107270	XXL	8600	k						
Skap 3											112815	Knatter	3600	d	
Catelli	111918	Dent Salt	6800	d											
Dynaflow	112815	Knatter	4300	k											

Figure 2.5: Brynild's drying cabinet schedule of Week 40 2020

To check if the schedule is feasible, the scheduler makes a visualization like we present in Figure 2.6.

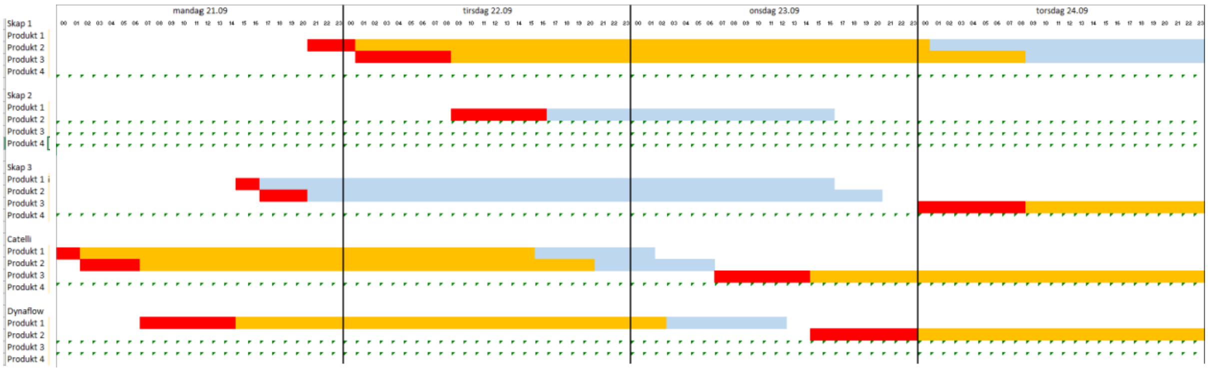

Figure 2.6: Visualization of a 3-shift drying schedule

Red presents the moulding, orange the heating, and blue the cooling. As examples on how to read Figure 2.6: in Skap 1 (the most upper one) 2 intermediates are dried at the same time next, no moulding is done between Tuesday afternoon and Wednesday end of the morning as all drying cabinets are occupied.

Brynild's scheduling procedure presents various challenges and restrictions: different shifts, changeovers, drying capacity, and shelf life. We elaborate these challenges in Section 2.2.1 to Section 2.2.4 respectively.

2.2.1 2-shifts and 3-shifts systems

Brynild uses various shift systems. A normal week consists of a 2-shift system or a 3-shift system. In a 2-shift system, there are 9 production shifts per week. Brynild uses the early and late shifts of the weekdays for production, except the late shift on Friday. There are 8 production hours per shift, which make up 72 production hours per week. In a 3-shift system, there are 14 production shifts per week. Brynild uses the early, late and night shifts of the weekdays for production, except for the late shift and the night shift on Friday, instead the production starts at Sunday night. In a 3-shift system, there are 8 production hours per shift as well, which make up 112 hours of production per week. For both types of shift-weeks there is a possibility for extension with 2 additional shifts. The additional shifts can be a late Friday shift and/or an early Saturday shift, see Figure 2.7.

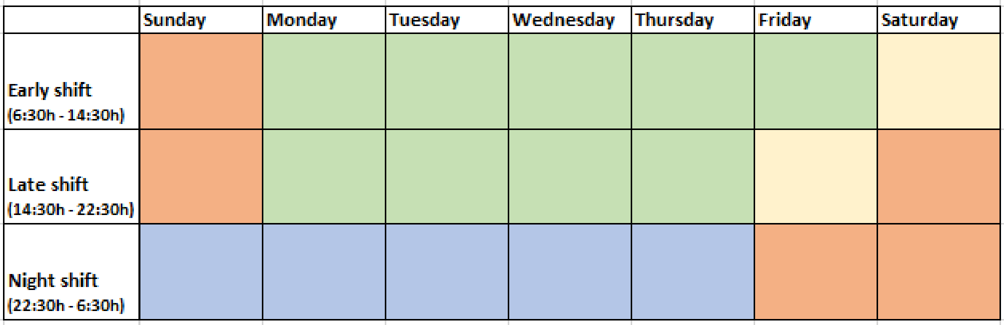

Figure 2.7: Brynild's different shift schedules

The green presents the 9 available shifts with a normal 2-shift system. Blue are the additional shifts that occur in a the 3-shift schedule. The yellow presents the shifts that Brynild can additionally schedule as overtime when 9 or 14 shifts are not enough. Orange indicates the time that Brynild cannot use for production. However, the drying continues through the whole week including during the orange times.

Overtime is paid 200% of the salary, which makes it far from ideal to run a system that always indicates to use these 2 additional shifts. One overtime is more expensive than one night shift, where manning is paid 130%. Another alternative for overtime is from 22:30h to 00:00h. Brynild can make this decision even on the same day, as the manning is normally more than willing to work these additional hours. Brynild decides at least 2 full weeks in advance whether they use a 2-shift or a 3-shift schedule.

2.2.2 Set-up times for changeovers

The set-up times for the changeovers in Støperi 1 are sequence dependent. The changeovers always happen within a shift or in overtime at the cost of available production time. The set-up time for a changeover depends on the tool that is required for the intermediate. Intermediates require different tools on Støperi 1. We give an overview of which intermediate requires what tool in Appendix A. When Brynild produces the same intermediate; no set-up time is needed. When the intermediates require the same tools only cleaning is necessary for which Brynild calculates 90 minutes of set-up time. When different tools are required, the set-up time is 120 minutes. These set-up times are upper bounds as we do not know beforehand, which employee does the changeover.

2.2.3 Drying area

The schedule in Figure 2.5 only shows the drying cabinets of Støperi 1. This is the main focus of the scheduler. Only when the scheduler can find some spare time, he focuses on Drage Sukker and Godteri Pakking as well. Otherwise, the operators of these lines have the job to process whatever is on the work floor. The reason that the scheduling mainly focuses on the drying area is due to several constraints for the drying cabinets. Some of these constraints we already mention in Section 2.1.2; limited capacity, intermediates drying together, different drying times per intermediate, and 1 family that is always dried in the new drying cabinets. In this section, we discuss the additional constraints for the drying cabinets and summarize all the constraints in Table 2.2.

The moulding trays circle from the drying cabinet to the moulding machine and back to the same drying cabinet. Familie K is dried at room temperature. The intermediates of this family cannot dry consecutively at the same drying cabinet more than 2 times, as the trays otherwise become too wet, because the temperature is too low to dry them sufficiently.

Another constraint is that Brynild never schedules the drying cabinets using full capacity. All drying cabinets must include at least 2 empty pallets. They use these empty pallets for flexibility. The cooking quantity is variable, and it is not exactly known how much more or less the operators produce in practice. The 2 empty pallets serve as buffer in case of overproduction. Using the 2 empty pallets, Brynild almost never has to discard confectionery mass. In this research we assume that this 2-pallet strategy is a good way to deal with the variability in quantity of the cooking process.

The final constraint regards the unloading of the drying cabinets. The intermediates that dry at room temperature need to be unloaded within 48 hours after finishing. The heated intermediates can stay in the drying cabinet for maximum 7 days. These 7 days include the drying time.

Table 2.2: Constraints drying area

Constraints	Clarification
Capacity	3 drying cabinets with 4 lanes and 2 drying areas with 2 times 2 lanes
Product families	Some intermediates within a product family are able to dry together
Throughput time	Old and new drying cabinets with different throughput time, which results in different drying times for the same intermediate
Familie D	One product family that can only dry in the new drying cabinets
Trays room temp.	Not more than 2 times in a row can room temperature intermediates dry in the same drying cabinet
Moulding and demoulding	A drying cabinet has to be available while moulding and demoulding
Sequencing	The sequencing with different set-up times affects the drying cabinet utilization
Two empty pallets	At least 2 empty pallets per drying cabinet scheduled for flexibility
Unloading room temp.	Intermediates that dry at room temperature need to be unloaded within 48 hours after completion
Unloading heated	Intermediates that dry with heat need to be unloaded within 7 days after entering the drying area

2.2.4 Packaging line and Shelf life

The confectionary products have shelf lives that vary from 5 up to 24 months. The retailer chains demand that there is a remaining shelf life of at least 100 days before shipment. The shelf life starts after packaging. Brynild prints the best before date on the package once the product is packaged, regardless of when the intermediate is produced. Two packaging lines, Bosch and HGD share the same subsequential line, called Roverma. Roverma is a wraparound line, which packs the consumer package into a distribution package. A consequence of sharing the same wraparound line is that Bosch and HGD cannot run at the same time.

2.3 CURRENT PERFORMANCE OF THE SCHEDULING PROCESS

There is no data on the realized performance of the schedules. Therefore, we evaluate Brynild's schedules instead. The data we use in this section we obtain from 8 consecutive schedules from 2020, like Figure 2.5. We assume that the drying area is empty at the beginning of the year due to a two-week Christmas break. The data we obtain show several flaws and incorrections:

- For 30% of the intermediates, the occupation time scheduled is less than the time that is required for occupation; filling and drying of the drying cabinets. For 8% of the intermediates, the occupation time scheduled is even less time than the minimum drying time;
- Around 4% of the intermediates are kept in the drying area longer than 7 days;
- Around 1% of the intermediates have a higher quantity when they exit the drying cabinet than when they enter;
- Around 3% of the intermediates that exit the drying area never entered;
- Around 3% of the intermediates are switched from drying cabinet, they exit a different cabinet than they enter.

We clean the data discrepancies with the assumption; what exits the drying cabinets is the most up to date data. Therefore, we make the following changes:

- We do not change the drying time to make the schedule feasible. However, in determining the utilization, we use an upper bound of 100% for the drying time. Therefore, the drying cabinet cannot be utilized more than 100%;
- We do not change anything in the schedule where intermediates are kept in the drying cabinet longer than 7 days, besides that this should not happen, the data is not contaminated by this discrepancy;
- When the quantity is higher when the intermediate exits the drying area, we assume that more Bryniled produces more quantity than thought beforehand. Therefore, we change the quantity that enters the drying cabinet to the same quantity that exits the drying cabinet;
- When an intermediate never enters a drying cabinet, but does exits one, we add this data. The schedule does mention when Brynild empties this drying cabinet before this specific intermediate. Thus, we add in the schedule that the specific intermediate enters the drying cabinet at this time;
- When the intermediates are switched from drying cabinets when they exit, we change the drying cabinet that they enter. So, what comes out of the drying cabinet is leading.

After cleaning this data, we research the different KPIs, which we mention in Chapter 1. In Section 2.3.1 we review the percentage of due dates reached, in Section 2.3.2 we discuss the occupation of the bottleneck, in Section 2.3.3 we calculate the utilization of the bottleneck, in Section 2.3.4 we compare the utilization of the bottleneck to its occupation, in Section 2.3.5 we present the make span of the production weeks, and in Section 2.3.6 we analyse the changeovers and their set-up times.

2.3.1 Percentage of due dates reached

To calculate, which percentage of the due dates Brynild reaches, we set the order due date indicated by the planning as a benchmark. If this order due date is not reached in terms of enough kg produced on time, we view this production order as unfulfilled. It does not matter if this order is unfulfilled by 1 kg or 1000 kg, or by 1 day or 100 days, as we view the fulfilment as a binary value. We assume that these production orders are backordered and therefore need to be fulfilled before the next production order of the same intermediate can be fulfilled. In the situation we describe, 64% of the order due dates are unmet.

Figure 2.8 presents the unmet percentage of due dates per week, the number of orders per week, and their trendlines.

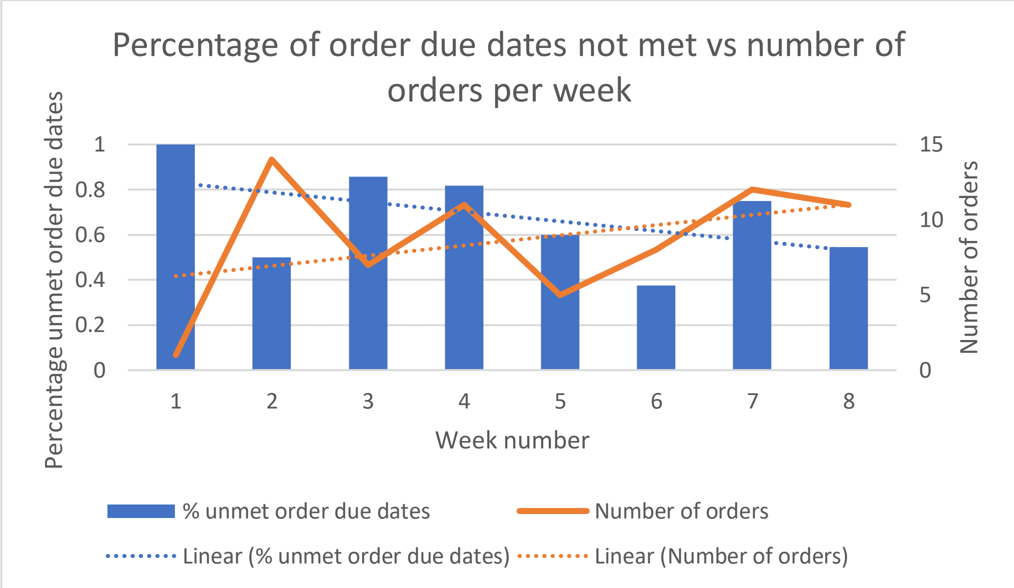

Figure 2.8: Number of unmet orders vs the number of order due dates per week

For the (linear) trendlines in Figure 2.8, we observe a descending trend in the percentage of unmet production orders per week and an ascending trend in the number of production orders per week.

We justify the observation of a descending trend of the percentage of unmet production orders per week with: needing time to re-stock inventory. In the Christmas holiday, Brynild closes, so there is no production, however there are still demand requests. Therefore, after the holiday break the inventory is lower than usual. This lower inventory makes it harder to meet the production order due dates. Likewise, the trend of number of orders per week is ascending. Brynild does not accept all production orders, when they know with certainty that they cannot be met. We explain both trendlines by the lower inventory after the Christmas holiday.

Week 2 however, does not fit in with the explanation, we explain the lower percentage of unmet order due dates by having 2 additional production days in Week 1. Brynild starts in Week 1 with producing the order due dates for Week 2, instead of the order for Week 1. This choice results in no order due dates met in Week 1, and more due dates met than we expect in Week 2. The high number of production orders in Week 2 is because Brynild knows beforehand that more production capacity is available than normally, as the order in Week 1 would require only 25% of the capacity of Week 1. When we discard Week 1 and Week 2 from this evaluation, we still observe that the trend lines follow the exact same trends as in Figure 2.8.

2.3.2 Occupation of the indicated bottleneck

After interviews with Brynild employees and several data analysis, we indicate the drying area as the bottleneck. We review the occupations of the cabinets separately, see Figure 2.9. With the term occupation we mean the time that intermediates are physically in the drying cabinet and the time that is required for moulding the intermediates.

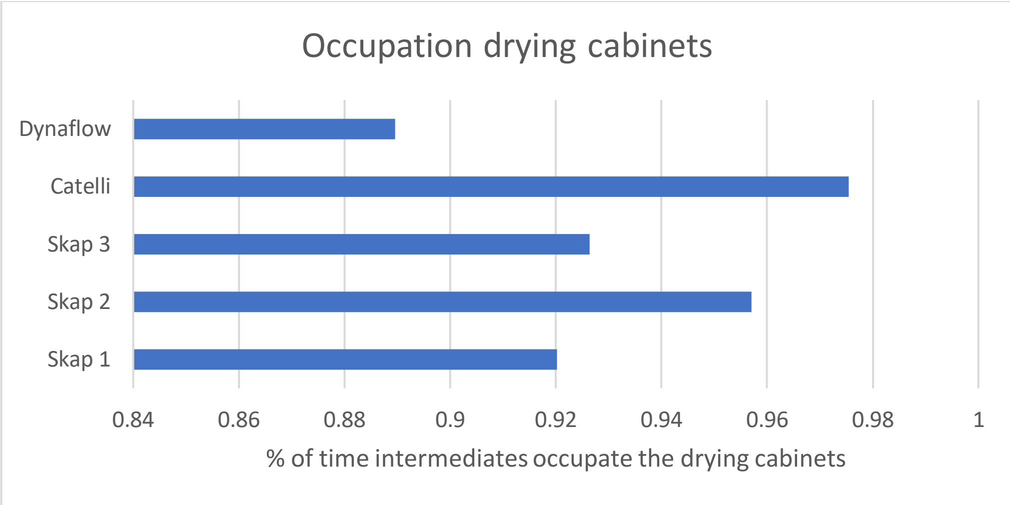

Figure 2.9: Occupation of the drying cabinets in %

Dynaflow has the lowest occupancy rate, namely 88,96%. We expect this high occupation percentages for two reasons. Firstly, Brynild's mention of the drying cabinets being almost always occupied. Secondly, Brynild generally only empties the drying cabinets when another intermediate enters the drying cabinet. Resulting in the drying cabinet being occupied for most of the time and used as storage area; waiting for Brynild to empty when another intermediate is ready to enter the drying cabinet.

2.3.3 Utilization of bottleneck

In Section 2.3.2 we observe high occupation percentages. In this section we calculate the utilization. When intermediates are in the drying cabinet, however the cabinet is not actually drying or cooling, the cabinet is not utilized. The cabinet is therefore occupied, but not utilized. Next to this, when the cabinet is drying intermediates, however half empty, the utilization is only half as well. We calculate the utilization for the 8 weeks as follows:

Utl = Utilized hours and space per intermediate per drying cabinet, see formula [2.1.1]
$\#L$ = Number of lanes used per drying cycle
T = Time needed for the drying and the moulding per drying cycle

$$Utl = \#L * T \qquad [2.1.1]$$

$TUtI$ = The total utilized hours per drying cabinet over the past 8 weeks, see formula [2.1.2]

$$TUtI = \sum Utl \qquad [2.1.2]$$

We calculate the $TUtI$ for every drying cabinet separately.

$TotTS$ = Total time and space available in the past 8 weeks, see formula [2.1.3]
$\#Hp8W$ = The number of hours per 8 weeks, which is 1304
$\#LpD$ = The number of lanes per drying cabinet, which is 4

$$TotTS = \#Hp8W * \#LpD = 5216 \qquad [2.1.3]$$

Thus, there are 5216 lane hours available per cabinet in the 8 weeks. Drying also occurs during the weekends and nights while there is no production during these times. We calculate the number of hours per 8 weeks by obtaining the difference between the first day and the last day of production and multiply the day difference by 24 hours. We calculate the total utilization per drying cabinet in the 8 weeks in formula [2.1.4].

$$Utilization\ per\ drying\ cabinet = \frac{TUtI}{TotTs} \qquad [2.1.4]$$

We present the results from formula [2.1.4] and the occupation we calculate in Section 2.3.2 in Figure 2.10.

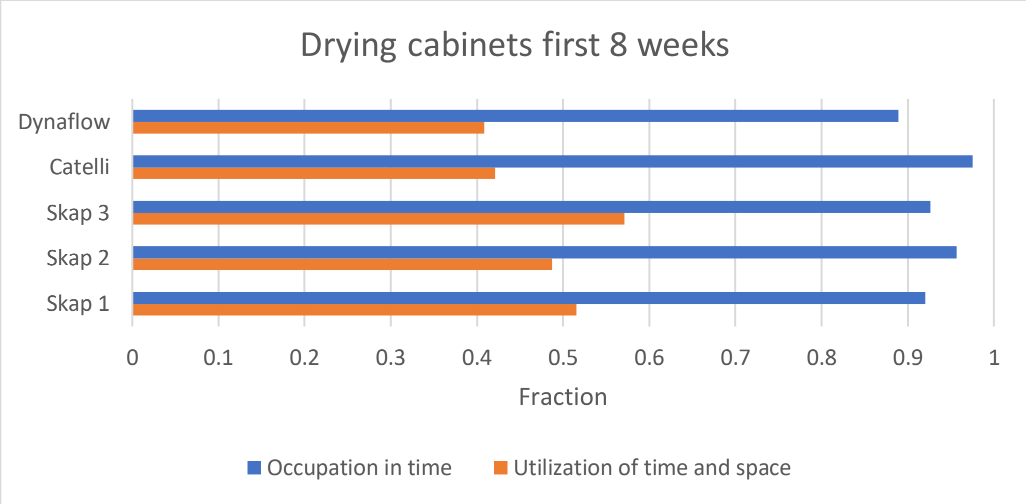

Figure 2.10: Utilization and Occupation of the drying cabinets

In Figure 2.10 we compare the utilization with the occupation from in Figure 2.9. The utilization is between 40% and 57%. When compare this percentage to the percentage of occupation, we observe that the utilization percentage is substantially lower. We expect that the utilization percentage and the occupation percentage differ somewhat, as the drying cabinets are not scheduled to be completely full. Next to that, drying finishes during the weekends or during the night, when the drying cabinet cannot be emptied immediately. However, the percentages should not differ this much. The occupancy percentage is, in comparison to the utilization percentage, approximately twice as high. Therefore, in theory on the base of Figure 2.10, using a 3-shift schedule the drying utilization could be almost doubled.

2.3.4 Utilization of bottleneck in comparison to the occupation

We also calculate the utilization in another way. Namely, not in comparison with the time that the drying cabinets could dry in the 8 weeks, but in comparison with the time the intermediates are occupying the drying cabinet. In this case, we calculate the effective drying time by dividing the time the intermediates require for moulding and drying by the time the intermediates are occupying the drying cabinet. We calculate the effective capacity by dividing the occupied number of pallets by the number of pallets available in that drying section, which is 36 or 72 depending on the drying cabinet. Using this calculation, we obtain the percentage of utilization while the drying cabinets are occupied. Thus, the percentage that the drying cabinets, when in use, are in use effectively. The higher percentage of the utilization in comparison with the occupation, the better. Figure 2.11 shows the effective drying time, effective use of the capacity, and utilization in comparison with the occupation per drying cabinet.

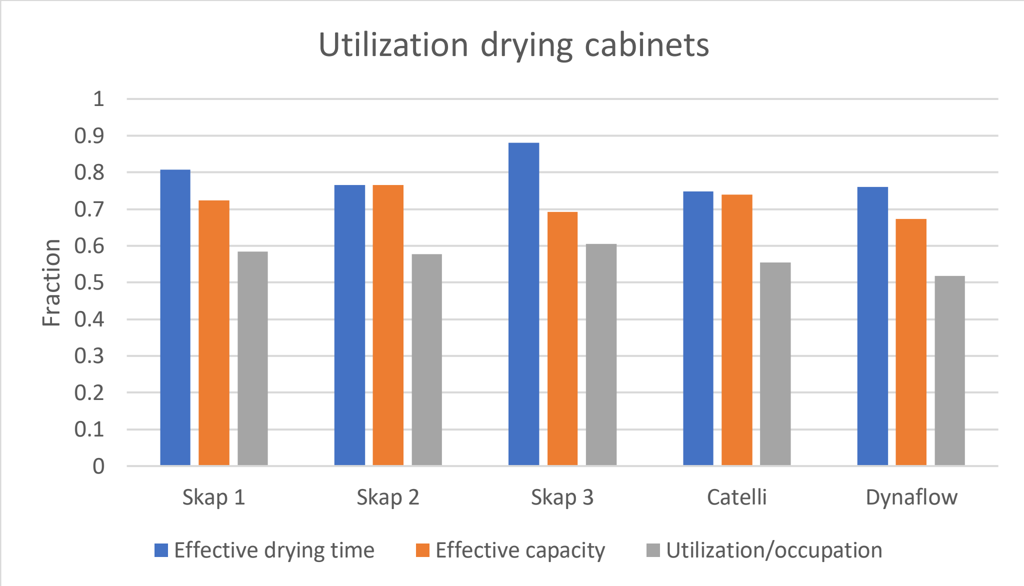

Figure 2.11: Different types of utilization per drying cabinet

Skap 3 has the highest effective drying time. Highest effective drying time means that Brynild uses this drying cabinet the least as storage area. We explain the relatively high effective drying time as 30% of intermediates that are scheduled use less time than what is needed in theory. This 30% of the intermediates have an effective drying time of 100%. Without these intermediates, the average effective drying time for the 8 weeks is 71% instead of 79%. The effective capacity utilization mostly dependents on the planning as the planning decides how much kg of each intermediate needs to be produced in a week. The schedule can influence the effective capacity utilization by scheduling more intermediates to dry together whenever possible.

2.3.5 Make span of the production week

With make span of the week, we intend the time it requires for all the intermediates that Brynild produces in a certain week to finish drying. For the make span, we calculate the number of hours between the start of a week and the time the last drying cabinet finishes drying, see Figure 2.12.

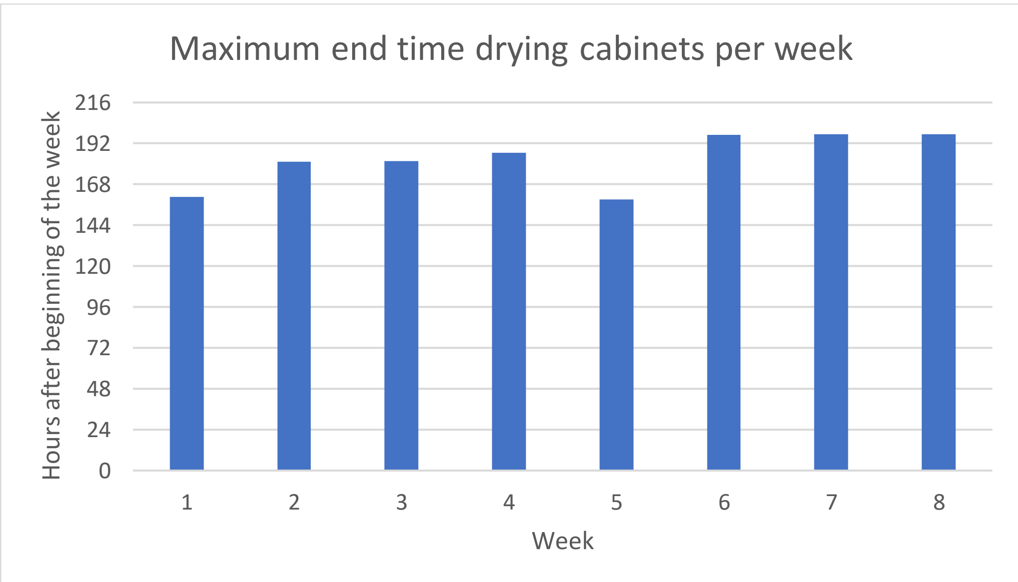

Figure 2.12: Make span per week

Weeks 6 to Week 8 have the longest make span. This is because in these weeks Brynild uses a 2-shift + Saturday schedule. In Week 1 and in Week 5 all drying cabinets finish on Monday. In none of the weeks all drying cabinets finish before the start of the upcoming week. This results in a lot of flexibility constraints for the next week, which results in a higher probability of obtaining a sub-optimal schedule. Figure 2.12 does not give much more information at this moment, however we use this figure as a benchmark for the new schedules.

2.3.6 Changeover set-up times

The average set-up time per changeover is 88 minutes and in total for the 8 weeks 7440 minutes. Figure 2.13 shows the number of times a type of changeover occurs in the 8 weeks.

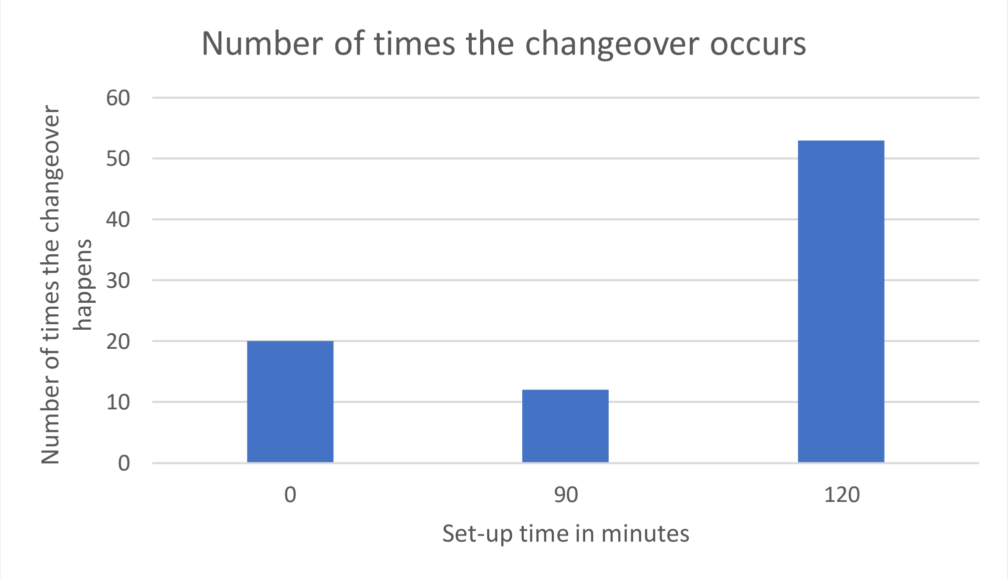

Figure 2.13: Number of times certain set-up times are required

Most changeovers need 120 minutes for set-up, see Figure 2.13. Brynild does try to schedule the same intermediates consecutively, however they do not focus on the same tool changeovers (90 minutes). We also evaluate the set-up times per week and the number of changeovers per week in Figure 2.14.

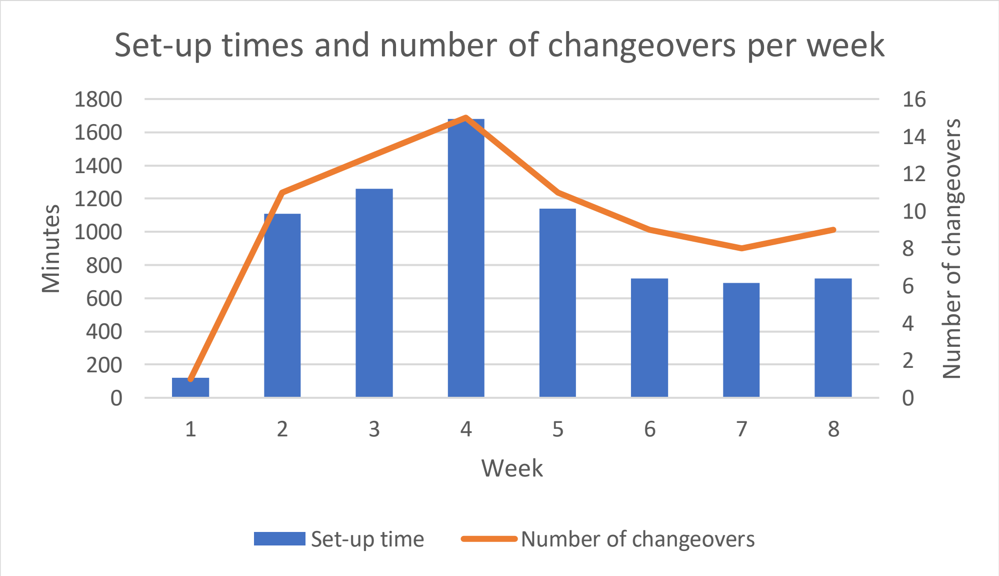

Figure 2.14: Total set-up time and number of changeovers per week

Figure 2.14 shows as we expect that there is a relation between the number of changeovers and the set-up time per week. The more changeovers, the more set-up time. The average set-up time per changeover per week is the highest for Week 4. We explain this by the high number of changeovers. Next, the higher the number of changeovers, the more difficult it is to have an overall overview of the schedule. With less changeovers like in weeks 6 to 8, it is easier for the scheduler to focus on the changeovers and not only on the feasibility of the schedule.

2.4 Conclusion

The confectionary production line consists of 3 parts. We give a detailed visualization in Figure 2.15.

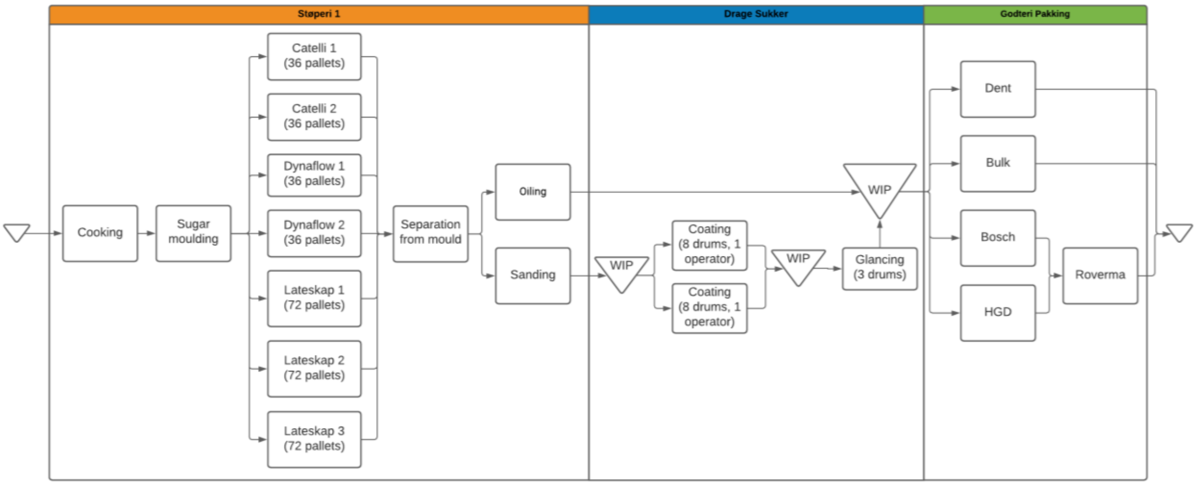

Figure 2.15: Brynild's confectionary production line in detail

We indicate the drying area as the biggest bottleneck of the entire production process. Therefore, the scheduler already focuses first on Støperi 1 and when more time is available the focus is on Drage Sukker and Godteri Pakking. The scheduler creates the production schedules of Støperi 1 manually and experience-based, taking into account the following two rules:

1. The production orders need to fit in the drying cabinets;
2. Try to use the new drying cabinets as much as possible.

We formulate the scheduling problem of Støperi 1 as: "*a problem with non-identical parallel machines, sequence depending set-up times, production time windows and without buffers*". The drying cabinets are the non-identical parallel machines as we discuss in Section 2.2. In Section 2.2 we discuss that the sequence dependent set-up times are zero when the same intermediate is produced, 90 minutes when the intermediates require the same tools, and 120 minutes when the intermediates require different tools. Last, we discuss that Brynild uses 2-shift schedules and 3-shift schedules therefore, production time windows are required. In Section 2.1 we mention that in Støperi 1 there is a continuous process starting from cooking till the intermediates entering the drying cabinets. This results in no buffer points between cooking and the drying area.

During the data analyses of the current performance of the schedule we observe that the bottleneck, the drying area, is almost always occupied, however not fully utilized. On average 50% of the time that the drying cabinets are occupied, they are used as storage area rather than actually being utilized. The two scheduling rules do not involve set-up times. We do observe however, that the scheduler does take these set-up times into account when the problem is small and does not require much changeovers. However, when the problem becomes bigger and more difficult to solve, the scheduler shifts his focus from the changeovers to solely creating a feasible schedule.

To identify and find possible solutions for this scheduling problem, we conduct a literature study in Chapter 3.

3 LITERATURE STUDY

This chapter answers the research question: '*Which methods are described in literature regarding the scheduling of a production line similar to Brynild's?*'. In this chapter we provide insight into alternative solutions for scheduling a production line that we can propose to Brynild. In general, scheduling problems are complex and dynamic, therefore they can be difficult to solve to optimality. The complexities arise from the many interrelationships that exist in scheduling problems (Williams, Pitts, & Kamery, 2004). In Chapter 2 we indicate that Brynild has a scheduling problem containing non-identical parallel machines (drying cabinets), sequence depending set-up times, production time windows and no buffer points. Therefore, the two most important decisions for Brynild's scheduling problem are the sequencing/start time of the intermediates and the assignment of intermediates to drying cabinets. In Section 3.1 we research different scheduling problems, with their roots in sequencing. In Section 3.2 we research whether we can identify Brynild's problem as an assignment problem (as well), because the assignment of the intermediates to the drying cabinets is an important part of Brynild's problem. In Section 3.3 we determine to what type of problem we can define Brynild's problem using information we obtain in Section 3.1 and Section 3.2. Section 3.4 contains a more elaborate mathematical description of the problem we select in Section 3.3. In Section 3.5 we analyse different solution approaches for the problem. Lastly, in Section 3.6 we conclude this chapter by giving an overview of our literature research and determine the solution approach we use.

3.1 THE DIFFERENT SCHEDULING METHODS

We start in Section 3.1.1 by discussing scheduling in general to narrow down the options in scheduling problems. After narrowing down we focus on the remaining scheduling problem, we start reviewing job shop scheduling in Section 3.1.2. Thereafter, in Section 3.1.3 we analyse a flexible job shop problem that includes 2 stages and parallel machines. In Section 3.1.4 we study the hybrid flow shop problem, which is somewhat more restricted than the flexible job shop problem.

3.1.1 Scheduling in general

There are three dimensions to classify production schedules (Graves, 1981):

1. Requirements generation;
2. Processing complexity;
3. Scheduling criteria.

The first dimension makes a distinction in terms of an open shop versus a closed shop. In an open shop the production orders are requested by customers and no inventory is stocked. In a closed shop the customer requests are serviced from inventory. We also know this as customer order decoupling point (CODP), see Figure 3.1.

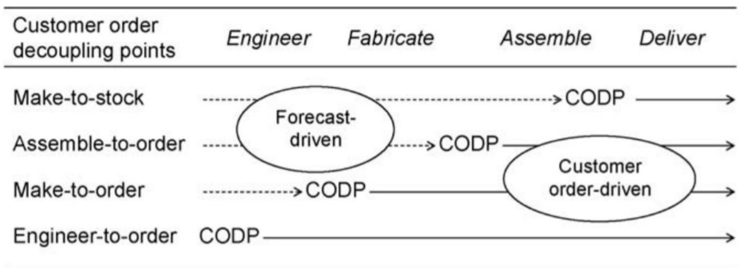

Figure 3.1: CODP options for different production lines

Brynild's production tasks are a result of inventory replenishment decisions. Therefore, Brynild is a closed shop scheduling problem and has a Make-to-Stock production line.

The second dimension is primarily concerned with the number of processing steps associated with each production task. A common breakdown is the following:

- One-stage, one-processor;
- One-stage, parallel processors;
- Multistage, flow shop;
- Multistage, job shop.

In multistage problems each task requires processing at different machines, where typically a strict precedence ordering for the machines is in state. The flow shop problem assumes that the tasks are to be processed on the same set of machines with an identical precedence. The job shop problem (JSP) is the most general production scheduling problem. In this classification, there are no restrictions on the processing steps and alternative routing is allowed. We describe Brynild's problem as a two-stage problem with a single machine in the first stage and parallel machines (drying cabinets) in the second stage. Brynild's problem is not a typical two-stage problem, because there are no buffer points between the first stage and the second stage, as this is a continuous process. Both the flow shop and the job shop can have single and parallel machines in different stages.

The third dimension indicates the measures that are used to evaluate the schedules. Two types of objective functions are schedule cost and schedule performance. The cost includes fixed costs, associated with set-ups, variable production and overtime costs, inventory costs, shortage costs, and expediting costs. The performance can be measured in many ways. Common measures are the utilization level of the resources, tardiness, and flow time. For Brynild's problem the performance is more important than the costs as the planning already considers the costs.

In Section 3.2 we research the processing complexity further by analysing the job shop problem and the flow shop problem.

3.1.2 Job shop scheduling

The job shop problem is first introduced by Muth & Thompson (1963). A simple job shop problem contains n jobs that need to be processed on m machines; the machines can process 1 job at the time. The processing of a job on a machine is called an operation and cannot be interrupted. The operation requires a given time on that specific machine. Each job consists of 1 or more operations, which need(s) to be completed in a specified sequence, forming a chain of operations. The sequencing of a job shop is allocating time on a machine for the operations, the sequence itself should be feasible and optimize the objective, which is classically minimize the make span.

The shifting bottleneck procedure (Adams, Balas, & Zawack, 1988) is a well-known (approximation) method of the JSP. This procedure gives a mathematical impression of the JSP and how sequencing is mathematically formulated in general.

$$
\begin{array}{llll}
Minimize & t_n & & (3.3.1) \\
Subtject\ to & t_j - t_i \geq d_i & \forall (i,j) \in A & (3.3.2) \\
 & t_i \geq 0 & \forall i \in N & (3.3.3) \\
 & t_j - t_i \geq d_i \; V \; t_i - t_j \geq d_j & \forall (i,j) \in E_k, k \in M & (3.3.4)
\end{array}
$$

In this method $N = \{0,1, \dots, n)$ denotes the set of operations (jobs or tasks), where 0 and n are the dummy operations for start and finish. M is the set of machines, A the set of precedence relations of the operations, and E_k the set of operations to be scheduled on machine k. The processing time of operation i is denoted by d_i and t_i refers to the start time of operation i.

The objective (3.3.1) refers to the make span, and minimizes the latest start time of the last operation, taking the following constraints into account:

- Constraints (3.3.2) ensure that precedence relationships are considered;
- Constraints (3.3.3) make sure that the starting time of a job cannot become negative;
- Constraints (3.3.4) ensure that only one operation is processed on a machine at the same time.

There are different job shop problems that consider resources and/or multiple plants (Zhang, Ding, Zou, Qin, & Fu, 2017). These JSPs with resources and/or multiple plants are out of scope for this research as they do not apply to the Brynild case. We do consider multiple stages; therefore, we analyse the flexible job shop problem (FJSP) in Section 3.2.2. The FJSP is an adaption on the JSP that includes the multi-purpose machine extension.

3.1.3 Flexible Job Shop Problem

The FJSP includes multi-purpose machine extension, which is a workplace with multiple machines and the operation can be processed on any among a set of available machines. Within this workplace the operation needs to be appointed to one of the machines. The operations might have different processing times on different machines (Pezzella, Morganti, & Ciaschetti, 2007). The FJSP introduces a further decision level besides the sequencing; namely the assignment of operations to suitable machines, also called routing. The FJSP can be categorized into two subproblems (Chaudhry & Khan, 2016):

1. A routing subproblem, where a suitable machine, among the available ones, needs to be selected to process the operation;
2. A scheduling subproblem, where a feasible schedule is obtained by sequencing the assigned operations, thus a sequencing problem.

There are two types of approaches to tackle this problem: hierarchical approaches and integrated approaches. In hierarchical approaches the assignment of operations to machines and the sequencing of the operations are treated separately. Thus, the assignment and sequencing are considered independently, this reduces the complexity of the problem (Xia & Wu, 2005). The first time this decomposition of the FJSP is used, is by Brandimarte (1993). The routing problem is solved using existing rules and thereafter the focus is on the sequencing problem. In integrated approaches, there is no differentiation between assignment and sequencing (Xia & Wu, 2005). The first integrated approach is presented by Dauzère-Pérès, Roux, & Lasserre (1998) where the neighbourhood structure made no distinction between reassigning and resequencing an operation.

There is a wide ranch of different objectives for the FJSP, some examples are: make span, mean completion time, total tardiness, maximum lateness, total workload of machines and critical machine workload. Often multi-objective optimization algorithms are used to achieve the minimum make span and the minimum cost with a balanced workload on machines. Thereafter the operation sequence is obtained (Zhang, Ding, Zou, Qin, & Fu, 2017). The Pareto approach provides an alternative to the multi-objective optimization. Different solutions are compared based on the Pareto dominance relation. This results in a set of multiple solutions that are all nondominated, meaning that none of the objectives can be improved without making one of the other objectives worse (Chiang & Lin, 2013).

3.1.4 Hybrid flow shop

In hybrid flow shops (HFS), n jobs need to be processed in a series of m stages. The sequence of the jobs needs to be determined in every stage to optimize a certain objective. The characteristics of the HFS are the following (Ruiz & Vázquez-Rodríquez, 2009):

1. The number of processing stages m is at least 2;
2. Each stage k has $M^{(k)} \geq 1$ machines in parallel and in at least one of the stages $M^{(k)} > 1$;
3. All jobs are processed following the same production flow. A job might skip a stage provided it is processed in at least one of them;
4. Each job j requires a processing time p_{jk} in stage k.

In general, in HFSs (i) all jobs and machines are available at time 0, (ii) the parallel machines at a given stage are identical, (iii) each machine can only process 1 operation at the time and each job can be processed by only 1 machine at the time, (iv) set-up times are negligible, (v) pre-emption is not allowed, (vi) the capacity of buffers between stages is unlimited, and (vii) the data is deterministic and known in advance.

The solution to a two-stage hybrid flow shop scheduling problem requires two aspects: the sequencing of the jobs on both stages and the assignment of the jobs to various machines at each stage (Gupta, 1988). Next to that, Gupta (1988) proves that the HFS problem is NP-complete by proving Theorem 1.
Theorem 1: For max (m_1,m_2)>1, the two-stage, hybrid flowshop problem is NP-complete even if the number of the machines at one of the two stages is one.

3.2 Assignment problems

In Chapter 1 we conclude that the drying area is Brynild's bottleneck. The scheduling of the drying area consists of two elements, sequencing of the intermediates, and the assignment of intermediates to the drying cabinets. This is a complex problem and in general to overcome complexities that arise from many interrelationships it is suggested to concentrate solely on optimizing a defined segment (Williams, Pitts, & Kamery, 2004).

In Section 3.1 we discuss different scheduling problems where we solely consider sequencing (JSP) and later consider both sequencing and routing/assignment (FJSP & HFS). In this section, we start by solely considering the assignment of the intermediates to drying cabinets, what we identify as an assignment problem (AP). The classic version of the AP is referred to as linear sum assignment problem (LSAP). The objective of the problem is finding the lowest cost assignment between n tasks and n agents, with 1 task per agent and 1 agent per task. The costs $c_{i,j}$, indicates the costs of assigning agent i to task j (Burkard, Dell'Amico, & Martello, 2009). In Section 3.2.1 we discuss the generalized assignment problem, which presents the assignment part of Brynild's problem best. In section 3.2.2 we analyse whether it is possible to incorporate sequencing in the assignment problem.

3.2.1 Generalized Assignment Problem

Many variations on the LSAP have been introduced over the years. The generalized assignment problem (GAP) is one of these variations and is first presented by Ross & Soland (1975). The goal for the GAP, is finding the lowest cost assignment between n tasks and m agents, where $n > m$, thus with multiple tasks per agent and 1 agent per task. The GAP is a NP-hard problem (Chakrabarty & Goel, 2010). Therefore, a GAP that determines the existence of a feasible solution is NP-complete and small problem instances can be solved within reasonable time (Martello & Toth, 1990).

The mathematical formulation of the generalized assignment problem is:

$$
\begin{array}{llll}
Minimize & \sum_{i\in I}\sum_{j\in J} c_{ij}x_{ij} & & (3.1.1)\\
Subject\ to & \sum_{j\in J} r_{ij}x_{ij} \leq b_i & \forall i \in I & (3.1.2)\\
& \sum_{i\in I} x_{ij} = 1 & \forall j \in J & (3.1.3)\\
& x_{ij} \in \{0,1\} & \forall i \in I,\ \forall j \in J & (3.1.4)
\end{array}
$$

Here, $I = \{1,2,\dots,m\}$ is a set of agents (or locations) indices, $J = \{1,2,\dots,n\}$ is a set of task indices, c_{ij} is the cost incurred for assigning agent i is to task j, r_{ij} is the resource required by agent i to carry out task j, and b_i is the amount of resource available to agent i. The objective (3.2.1) is to minimize the total assignment costs, considering the following constraints:

- Constraints (3.1.2) ensure that the capacity of the agents are not exceeded;
- Constraints (3.1.3) make sure that every task is assigned to exactly 1 agent;
- Constraints (3.1.4) are the domain constraints, where $x_{ij} = 1$ when agent i is assigned to task j, 0 otherwise.

This problem presents Brynild's assignment best, because in general Brynild has to assign more intermediates (tasks) to fewer drying cabinets (agents). In Brynild's case, assigning an intermediate to a specific drying area incurs time instead of costs. The amount of resource is in our problem the capacity of each drying cabinet. This GAP is only one part of Brynild's problem. After a static assignment, the intermediates still need move to the second stage and arrive at the assigned cabinet. In Section 3.2.2 we study the generalized quadratic assignment problem (GQAO) that incorporates movement.

3.2.2 Generalized Quadratic Assignment Problem

The GQAP is introduced by Lee & Ma (2003). The GQAP optimally assigns m pieces of equipment to n locations, where $m > n$. Tasks are performing a specific sequence of operations and must move between equipment. Each piece of equipment must be assigned to exactly 1 location, and there can be multiple pieces of equipment at 1 location, as long as there is no violation of the space limitation. The objective is to find an assignment such that the total costs, consisting of assignment costs and transportation costs, is minimized.

The mathematical formulation for the GQAP is:

$$
\begin{array}{llll}
Minimize & \sum_{i=1}^{m}\sum_{k=1}^{n} c_{ik}x_{ik} + v\sum_{i=1}^{m}\sum_{j=1}^{m}\sum_{k=1}^{n}\sum_{h=1}^{n} q_{ij}d_{kh}x_{ik}x_{jh} & & (3.2.1)\\
Subject\ to & \sum_{i=1}^{m} s_i x_{ik} \leq S_k & \forall k & (3.2.2)\\
& \sum_{k=1}^{n} x_{ik} = 1 & \forall i & (3.2.3)\\
& x_{ik} \in \{0,1\} & \forall i \in I, \forall j \in K & (3.2.4)
\end{array}
$$

Here, $i, j = \{1,2,\dots,m\}$ is a set of equipment and $k, h = \{1,2,\dots,n\}$ is a set of locations. The following parameters and variable are indicated as:

- c_{ik} is the cost incurred for assigning equipment i is to location k;
- q_{ij} is the flow volume (volume of all tasks) from equipment i to equipment j;
- d_{kh} is the distance between location k and location h;
- v is the travel cost per unit distance and per unit flow volume;
- s_i is the space required by equipment i;
- S_k is the available space at location k.

The objective (3.2.1) is to minimize the total costs, in assignment and transportation costs taking the following constraints into account:

- Constraints (3.2.2) ensure that the space capacity of the location k is not exceeded;
- Constraints (3.2.3) make sure that every equipment is assigned to exactly 1 location;
- Constraints (3.1.4) are the domain constraints, where $x_{ik} = 1$ when equipment i is assigned to location k, 0 otherwise.

The GQAP has a form of transportation costs. In Brynild's case this could be the time incurred by transporting the cooked mass to the drying area. However, this problem does not consider the occupation of a drying cabinet that comes with previous arrived and assigned jobs. The occupied drying cabinets and sequencing of intermediates is timing related. The timing is important to obtain a feasible solution for Brynild's problem.

Besides the GAP and GQAP we review other assignment problems. For example, the process allocation problem (Chu, 1969), which is considered a special class of the GQAP and the quadratic multiple knapsack problem (QMKP) (Hiley & Julstrom, 2006). The QMKP states that each task is assigned to at most 1 agent (location), which in our case could lead to no assignment for certain intermediates. Looking from the assignment perspective to Brynild's problem, we do not result in in finding a complete problem that is fully relatable problem to Brynild's case.

3.3 Determination scheduling problem Brynild: HFS

We summarize the capabilities of the problems we review in Section 3.1 and Section 3.2 in Table 3.1.

Table 3.1: Capabilities of the problems we review

	JSP	**FJSP**	**HFS**	**GAP**	**GQAP**
Parallel machines	No	Yes	Yes	Yes	Yes
Non-identical parallel machines	No	Yes	No	Yes	Yes
Sequencing/Timing	Yes	Yes	Yes	No	No
2-stages	No	Yes	Yes	No	Yes
Zero buffer space	No	No	No	No	No

The constraints for a general FJSP have the most in common with Brynild's scheduling problem. However, the production line of Brynild is a flow shop. All intermediates follow the exact same order on the production line. Brynild's production line is a special HFS case, which is far from the general HFS as:

(i) Not all machines are available at time 0, as some drying cabinets can still be drying the intermediates from the previous week;
(ii) The parallel machines are not identical, they can have different throughput times and capacity constraints;
(iii) Each machine can in some cases process more than 1 operation at the time, however each job can indeed only be processed by 1 machine at the time;
(iv) Set-up times are not neglected;
(v) Pre-emptions are indeed not allowed, same as for the general HFS;
(vi) Brynild has a two-stage process without buffer points between the single machine stage and multiple machine stage;
(vii) Same as for the general HFS the data is deterministic and known in advance.

An eighth point we like to mention is that Brynild has production time windows. This is not specifically mentioned for a general HFS, therefore we do not focus on the production time windows in the rest of the literature review. We do consider the production time windows in Chapter 4. Thus, Brynild's problem is a two-stage hybrid flow shop with non-identical parallel machines, sequence dependent set-up times, production time windows and without any buffer points between the 2 stages.

3.4 Two-stage HFS mathematical description

There are not many mathematical formulations regarding a two-stage hybrid flow shop that is similar to the problem we identify in Section 3.3. In this section we present a mathematical formulation for the HFS that is partially relatable to Brynild's problem. The HFS has a single machine in Stage 1 and identical parallel machines in Stage 2. Next, the HFS considers sequence dependent set-up times and unlimited buffer points (Lee, Hong, & Choi, 2015). In the formulation, job 0 is a dummy job to impose set-up time for the first job(s) and the processing time of job 0 is set to be zero.

$$Min\ z \tag{3.5.1}$$

$s.t.$

$$\sum_{j=1}^{N} x_{0j1} = 1 \tag{3.5.2}$$

$$\sum_{j=1}^{N} x_{0j2} = R \tag{3.5.3}$$

$$\sum_{j=1}^{N} x_{ijk} = 1, \qquad for\ i = 1,2,\dots,N,\ \ k = 1,2 \tag{3.5.4}$$

$$\sum_{i=0}^{N} x_{ijk} = 1, \qquad for\ j = 1,2,\dots,N,\ \ k = 1,2 \tag{3.5.5}$$

$$c_{j1} - c_{i1} + BigM(1 - x_{ij1}) \geq s_{ij} + p_{j1}, \qquad for\ i = 0,1,\dots,N,\ \ j = 1,2,\dots,N \tag{3.5.6}$$

$$c_{j2} - c_{i2} + BigM(1 - x_{ij2}) \geq p_{j2}, \qquad for\ i = 0,1,\dots,N,\ \ j = 1,2,\dots,N \tag{3.5.7}$$

$$c_{j2} - c_{j1} \geq p_{j2}, \qquad for\ j = 1,2,\dots,N \tag{3.5.8}$$

$$c_{jk} \geq c_{0k}, \qquad for\ j = 1,2,\dots,N,\ \ k = 1,2 \tag{3.5.9}$$

$$z \geq c_{j2}, \qquad for\ j = 1,2,\dots,N \tag{3.5.10}$$

$$x_{ijk} \in \{0,1\}, \qquad \forall i,j,k \tag{3.5.11}$$

$$x_{ijk} = 0, \qquad \forall i = j,\ \ k = 1,2 \tag{3.5.12}$$

$$c_{jk} \geq 0, \qquad for\ j = 0,1,\dots,N,\ \ k = 1,2 \tag{3.5.13}$$

Here, $i, j = \{1,2,\dots,N\}$ is a set of jobs and $k = \{1,2\}$ is the set of stages. The parameters are:

- s_{ij} is the set-up time between job i and job j;
- p_{jk} is the processing time of job j at stage k;
- R is the number of identical parallel machines;
- BigM is a very large value.

The variables in the formulation present:

- c_{jk} is the completion time of job j at stage k;
- x_{ijk} is 1 when job i is processed directly before job j at stage k, 0 otherwise.

The objective (3.5.1) is to minimize the make span, taking the following constraints into account:

- Constraints (3.5.2) and (3.5.3) ensure that the jobs are scheduled on a machine at Stage 1 and on R machines at Stage 2 by assigning job 0 to the R machines;
- Constraints (3.5.4) and (3.5.5) ensure that each job is sequenced immediately before and after only 1 job, respectively, on only 1 machine at each stage. The summation is over different ranges as job 0 does not have predecessor;
- Constraints (3.5.6) ensure that the completion time of job j at Stage 1 is greater than that of job i by at least the processing time of job j at Stage 1 plus the set-up time between jobs i and j, if job i is sequenced immediately before job j at Stage 1;
- Constraints (3.5.7) ensure that the completion time of job j at Stage 2 is greater than that of job i by at least the processing time of job j at Stage 2, if job i is sequenced immediately before job j on a machine at Stage 2;
- Constraints (3.5.8) ensure that the completion time of a job at Stage 2 is always greater than that of the same job at Stage 1 by at least the processing time of the job at Stage 2;
- Constraints (3.5.9) force job 0 to be sequenced first;

- Constraints (3.5.10) link the objective function and the decision variables;
- Constraints (3.5.11), (3.5.12) and (3.5.13) impose the boundary of the decision variables.

This formulation of a HFS is introduced by Lee, Hong & Choi (2015) to describe the scheduling problem in more detail. This mathematical model is not solved as it is a relatively large-sized problem, which needs to be solved within reasonable short time. There are differences between this two-stage HFS and Brynild's two-stage HFS. Firstly, Brynild's HFS includes unidentical parallel machines instead of identical parallel machines. Next, Brynild's problem needs to incorporate production time windows. Lastly, the HFS we analyse in this section has unlimited buffer capacity between the 2 stages, while Brynild has a continuous production line and therefore, zero buffer capacity between the 2 stages.

3.5 SOLUTION APPROACHES TWO-STAGE HFS

In this section we examine various approaches to solve a hybrid flow shop with:

- Non-identical parallel machines;
- Sequence depending set-up times;
- Zero buffer capacity.

We do not research the production time windows in this literature research as we mention in Section 3.3.

There is very limited literature on HFS problems without buffer capacity. We find 5 articles of which 3 articles consider a problem setting that is irrelevant for Brynild's problem (Wei, Wu, Jiang, & Cheng, 2019), (Zhang, Rong, & Liu, 2014) and (Chen, Pan, Zhang, Ding, & Li, 2019). Next, we analyse a promising abstract, however the remainder of this paper is written in Chinese, a language skill we do not possess (Zhang, Li, Storer, & Yan, 2013). Thus, there is 1 paper that partially represents Brynild's two-stage HFS without buffer capacity.

In Section 3.5.1 we research the two-stage HFS without buffer capacity and other HFS problems that include finite buffer capacity. We consider finite buffer capacity as well, to research the influence of the buffer points and what solution approaches are applied to those types of problems. In Section 3.5.2 we identify solution approaches for HFSs with sequence dependent set-up times and in Section 3.5.3 we discuss solution approaches for HFSs with non-identical parallel machines. We cannot find much literature on HFSs including a combination of 2 of Brynild's characteristics and no literature at all on the 3 characteristics combined. To broaden our knowledge, we research solutions for flexible job shop problems as well. A summary of this literature study is given in Appendix B. From this literature study we conclude that also in this area we are not able to find representing FJSPs without buffer point capacity and therefore no new insights in solution approaches for this type of problem.

3.5.1 HFS with finite intermediate buffers

The only paper that includes a two-stage hybrid flow shop problem without buffer points interesting for the Brynild case is applied to a hospital setting (Dekhici & Belkadi, 2010). The first stage in this paper are the operating rooms, and the second stage the beds to which the patient needs to be assigned. Dekhici & Belkadi propose a local search for constraint satisfaction in the initial feasible solution and a tabu search (TS) with a restricted neighbourhood system for a make span optimization. TS is a meta-heuristic based on the local search principle. TS begins from an initial solution and chooses at each iteration the best solution in the current neighbourhood, even when the best solution in the current neighbourhood does not improve the quality of the current best solution found. Therefore, TS can escape a local optimum. TS makes use of a tabu list that includes certain actions that are previously executed. The tabu list has a certain length and when the list is full, the principle first in first out occurs and one the previous action can be executed again. The efficiency of approaches based on local search, such as tabu search, depends heavily on the quality of the initial solution (Ennigrou & Ghédira, 2008). For the hospital problem the TS provides feasible schedules, however there is no comparison with other meta-heuristics to determine if TS creates good schedules (Dekhici & Belkadi, 2010).

We also research a hybrid flow shop problem with limited buffer capacity (Yaurima, Burtseva, & Tchernykh, 2009). This problem includes unrelated machines, sequence-dependent set-up time, availability constraints and limited buffers. In the paper they propose a genetic algorithm (GA) to minimize the make span. Many authors separate sequencing and assignment decisions in the HFS problem. However, in this paper they follow the way proposed by Ruiz and Maroto (2006), where the assignment of jobs to machines and the sequence in each stage is done by the evaluation function. The change Yaurima, Burtseve & Tchernykh (2009) make to this function is that a job is assigned to the machine that can finish the job at the earliest time at a given stage instead of assigning the job to the first machine available.

Another HFS scheduling problem with finite intermediate buffers is solved with a TS (Wang & Tang, 2009). The objective is to minimize the sum of weighted completion time of all jobs. A TS heuristic where a scatter search mechanism is incorporated is proposed. The TS has two neighbourhood structures; 1) reinsertion, and 2) swap, the scatter search is incorporated to improve the diversity of the search procedure. The solution denotation is a permutation of all jobs representing their processing order in the first stage and the greedy constructive procedure that obtains the corresponding complete schedule. The TS provides good solutions compared to the lower bound and outperforms the Lagrangian relaxation algorithm. The Lagrangian relaxation is a method that approximates the problem by incorporating adapted constraints as punishment in the objective, instead of using constraints.

3.5.2 HFS with sequence dependent set-up times

HFSs with sequence dependent set-up times (SDST) are commonly known as SDST hybrid flow shops. Very limited studies address this exact type of problem, as the set-up time is often ignored in literature or considered as a part of the processing time. Zandieh, Ghomi and Husseini (2006) compare an immune algorithm (IA) to a random key genetic algorithm (RKGA), to solve the SDST HFS. IA and RKGA are both generic algorithms. AI is based on the biological immune theory. An operator accomplishes immunity by two steps: 1) a vaccination and 2) an immune selection, of which the former is used for raising fitness and the latter is for preventing the deterioration. RKGA is based on an array of n random keys, where a random key is a real number randomly generated. The objective used in the paper is to minimize the make span and Zandieh, Ghomi and Husseini (2006) establish that in their paper the IA outperforms the RKGA.

Li (1997) considers a two-stage hybrid flow shop with a single machine in Stage 1 and multiple identical parallel machines at Stage 2. Li (1997) characterizes the flow shop by major and minor set-ups, part families, and batch production allowing to split jobs at Stage 2. The objective is to minimize the make span. Li (1997) develops 2 allocation policies: forward heuristic (FH) and backward heuristic (BH). The FH starts with sequencing of Stage 1 and allocates the jobs to machines afterwards. The advantage of the FH is that it is easy to understand and implement. The BH starts with allocating the jobs to machines at Stage 2 and sequences the jobs later. Li (1997) finds that the Backward Heuristic is in general superior to the Forward Heuristic.

Another HFS with sequence depending set-up times evaluates multiple meta-heuristics (Bozorgirad & Logendran, 2016). They comprehensively address the question of what type of meta-heuristic algorithm is most appropriate to solve these types of problems. The paper reviews tabu search, simulated annealing (SA) and the genetic algorithm. The objective of the scheduling problem is to minimize a linear combination of the total weighted completion time and the total weighted tardiness of the jobs. In general, the GA performs the best. However, when we only evaluate interesting instances to the Brynild case, we observe that in most instances the TS and/or the SA outperformance(s) or match(es) the results of the GA.

3.5.3 HFS with non-identical machines

HFS with non-identical machines is also referred to as HFS with unrelated parallel machines. The unrelated parallel machines have different run times for the same job due to different machine capabilities. We discuss an HFS with unrelated parallel machines where they consider pre-determined groups of jobs (Shahvari & Logendran, 2016). The objective is to minimize the weighted sum of the completion time and total tardiness. Shahvari & Logendran (2016) propose a meta-heuristic algorithm based on tabu search. The tabu search moves back and forth between batching and scheduling phases. Bozorgirad & Logendran (2013) also use a tabu search to solve an HFS with unrelated-parallel machines. In this paper the unrelated-parallel machines indicate that not all machines can process all jobs. This is relevant as some of Brynild's intermediates cannot be processed by every drying cabinet.

Aqil & Allali (2020) discuss a hybrid flow shop with unrelated parallel machines under constraints of sequence dependent set-up time. The objective is to minimize the total tardiness of all jobs. Two methods are proposed to solve the problem, the iterative local search (ILS) and the iterated greedy (IG) meta-heuristics. The ILS is based on two main phases, the first phase starts from an initial solution, it consists of exploring the neighbourhood by a perturbation procedure. The second phase, is allowing the local search to accept or reject the new solution according to a given criterion. The IG is based on two main steps, the first step consists of two phases, the destruction phase and the reconstruction phase. The second step is the exploration in local search, which is important for the search in the neighbourhood of the current solution. Aqil & Allali (2020) consider an initial solution set generated using priority rules. They classify the jobs in descending order based on the priority rules to obtain a sequence to number the jobs. The further scheduling of these solutions is ensured by the Nawaz-Enscore-Ham (NEH) algorithm and the greedy randomized adaptive search (GRASP) procedure. In the NEH two subsets of the first two jobs are created. Then an iteration from job 3 to n is conducted by inserting the new job in the different positions of the current built solution and remembering the sequence with the smallest tardiness. In the GRASP a job is selected from a restricted candidate list to be scheduled in the current sequence at an interval established by a margin of choice given by the objective function. Aqil & Allali (2020) find that the IG algorithm based on NEH initialization heuristic gives a good qualitive start solution that has a good convergence time to the best solution.

An improved artificial bee colony (IABC) algorithm for the HFS with unrelated parallel machines is also proposed (Li, Li, Gao, & Meng, 2020). The IABC adopts a greedy iterative strategy to generate high quality initial solutions. Next to that the IABC also adopts the advantages of simulated annealing and retention mechanism. In the paper they compare IABC to GA and IG, where IABC performs the best, however the iterative greedy algorithm is a very close second.

3.6 Conclusion

After an extensive literature review in Section 3.1 and Section 3.2, we come to the conclusion in Section 3.3 that the Støperi 1 line of Brynild is a two-stage hybrid flow shop with non-identical parallel machines, sequence dependent set-up times, production time windows, and without intermediate buffer points.

We study 6 different types of HFS and their solution approaches. We present the 6 different HFS in Table 3.2.

Table 3.2: Different HFS models

Model	Zero buffer points	Limited buffer capacity	Non-identical parallel machines	SDST
1	X			
2		X		
3			X	
4				X
5			X	X
6		X	X	X

In Table 3.3, we summarize the various solution approaches for the HFS problems from Table 3.2, and the source of the paper that we obtain the solution approach from.

Table 3.3: Heuristics used to solve the HFS problems and the source

	1	2	3	4	5	6
TS	(Dekhici & Belkadi, 2010)	(Wang & Tang, 2009)	(Shahvari & Logendran, 2016)	(Bozorgirad & Logendran, 2016)		
GA				(Bozorgirad & Logendran, 2016)		(Yaurima, Burtseva, & Tchernykh, 2009)
SA				(Bozorgirad & Logendran, 2016)		
IA				(Zandieh, Ghomi, & Husseini, 2006)		
IABC			(Li, Li, Gao, & Meng, 2020)			
RKGA				(Zandieh, Ghomi, & Husseini, 2006)		
ILS					(Aqil & Allali, 2020)	
IG			(Li, Li, Gao, & Meng, 2020)		(Aqil & Allali, 2020)	
FH				(Li S. , 1997)		
BH				(Li S. , 1997)		
Mathematical model				(Lee, Hong, & Choi, 2015)		

The HFS is a complex scheduling method. In general, to overcome the scheduling complexities that arise from many interrelationships, near optimal solutions are suggested, or it is suggested to concentrate solely on optimizing a defined segment within the total operational framework (Williams, Pitts, & Kamery, 2004). This is exactly what we see in our literature review, since all HFSs that we find are solved using (meta) heuristics to find a (near) optimal solution and the HFSs with limited buffer capacity optimize only one defined segment.

Using Table 3.3 we conclude that the tabu search is most frequently used to solve problems that are similar to Brynild's problem. From our research regarding solution approaches for FJSPs we conclude that TA and GA are the most common used as well. However, in Brynild's problem there is no buffer capacity between the 2 stages and production time windows are in place. Due to these two characteristics, is it difficult to obtain good solutions using an integrated meta-heuristic. When using an integrated meta-heuristic a lot of solutions the heuristic creates are infeasible. Some infeasible solutions can be useful to escape a local optimum, however too many infeasible solutions make it very difficult to obtain a (near) optimal solution

Another option to solve the problem is to separate the assignment and sequence and use a meta-heuristic on one part of the problem. Because when the assignment (sequence) changes the sequence (assignment) must most likely be adapted accordingly to create a feasible solution. Therefore, we can apply a meta-heuristic to the assignment (sequence) and enumerate the sequence (assignment) of the jobs to obtain the optimal sequence (assignment) for that assignment (sequence).

Luckily, the instances of Brynild's problem are likely small enough to solve within reasonable time using a Mixed Integer Linear Programming (MILP) model. In Chapter 4 we present our MILP model that represents Brynild's scheduling problem. To develop this Brynild's MILP we use the MILP from Section 3.4 as a guideline.

4 MILP DESCRIPTION OF BRYNILD'S PRODUCTION LINE AND 3 HEURISTICS

In this chapter we answer the following research question: '*How can we develop a scheduling model that improves the throughput of the production line?*'. As we mention in Chapter 3, we design a MILP including Brynild's uncommon characteristic; no buffer point capacity. To develop a scheduling model, we determine the scope and clarify our modelling assumptions in Section 4.1. In Section 4.2, we describe the mathematical model and the objective function to improve the throughput of the production line. In Section 4.3, we present 3 heuristics based on our MILP. Lastly, in Section 4.4 we conclude this chapter.

4.1 SCOPE AND ASSUMPTIONS OF THE MATHEMATICAL MODEL

Recall that first Brynild focuses solely on the Støperi 1 line. In our solution approach we focus only on Støperi 1 as well, as this line contains the drying area, the bottleneck. If we consider the entire production line, including the packaging line, the model becomes very large and complicated. Due to the buffer points between Støperi 1, Drage Sukker, and Godteri Pakking, we consider these 3 lines as 3 separate scheduling tasks, as we assume that there is infinite buffer capacity between the lines. In practice Brynild's buffer capacity is limited, however Brynild indicates that the specific production day within a week for each intermediate has little effect on the buffer capacity. For the capacity of the buffers, it is far more important in what week Brynild produces the intermediate. Therefore, we assume infinite buffer capacity between the 3 lines, as Brynild considers the buffer capacity while planning the production orders.

The main focus of scheduling Støperi 1 is the sequencing/start times of the production orders and the assignment of intermediates to the drying cabinets. The start times have a direct effect on how we can assign the jobs to drying cabinets. After demoulding, the intermediates are continuously transported to the oiling or sanding process in Støperi 1. Both the oiling and sanding process have enough production capacity to handle the output of the demoulding process, independent of the sequencing and assignment in the preceding processes. Therefore, we exclude the oiling and sanding process from our analysis. Concluding, we focus on the start times of the production orders and the assignment of intermediates to drying cabinets.

Given this focus, we model the scheduling problem as a two-stage hybrid flow shop, where the drying cabinets are non-identical parallel machines. We include sequence dependent set-up times, production time windows, and no buffer capacity between the 2 stages. The next step is to define how we construct and use the input data.

The complexity of the mathematical model depends on how much we pre-process the input data. During pre-processing, we make important decisions before solving the model. More pre-processing indicates a smaller solution space, and therefore the MILP is less complex and requires less computation time. We assume that the input data is structured in such a way that jobs consist of 1 or more intermediates, which can dry together and, in terms of quantity always fit in 1 drying cabinet. This assumption has the advantage that we do not need to include any constraints and variables for batching in the mathematical model, which results in a more comprehensible MILP. The advantage of pre-determining the batches is that the schedule we create connects better to practice as undesired outcome for quantity production and the sequence of intermediates drying together is already covered by the pre-processing. The disadvantage is the loss of flexibility. In Appendix C we have a detailed discussion concerning this assumption for the input data together with various alternatives. We conclude from this detailed discussion that we need to make another assumption as well. To obtain the input data structure we must assume that all drying cabinets have the same capacity as otherwise we cannot know the maximum size of a job. Therefore, we combine the 2 Catelli drying cabinets in the model and we also combine the 2 Dynaflow cabinets.

Next to the input data we make some other assumptions. We assume that the drying cabinets are available at the beginning of the week. This is a reasonable assumption as in a 2-shift schedule there are 64 hours available for drying in the weekend and in a 3-shift system there are 56 hours available. Table 4.1 presents the percentage of the intermediates that always finish drying during the weekend when Brynild produces them Friday afternoon.

Table 4.1: Percentage of intermediates to complete drying before the start of the new week

	New drying cabinets	Old drying cabinets
2-shift	97,4%	73,7%
3-shift	81,6%	68,4%

In the mathematical model we focus on the end time of the drying cabinets and not on the time when we empty the drying cabinets. Therefore, we do not consider the maximum time an intermediate is allowed to stay in the drying cabinet. This is reasonable as the drying cabinets can be emptied without having to re-fill the drying cabinet immediately. Emptying, without refilling does require more handling. However, the schedule should not depend on the fact that the drying cabinet is solely emptied when a new intermediate enters the drying cabinet, as this is not necessary for the production line.

Other additional assumptions are:

- All machines in the production line are available at all times; no breakdowns or maintenance;
- Jobs process always without error;
- We cannot interrupt job processing, so no pre-emption;
- We cannot process jobs on more than 1 machine simultaneously;
- A machine may only process 1 job at a time;
- All data is known deterministically when creating the schedule;
- We do not consider set-up time between weeks.

4.2 MATHEMATICAL MODEL

In this section we describe and explain the mathematical model of the two-stage HFS considering:

- No buffer points;
- Sequence dependent set-up times;
- Non-identical parallel machines;
- Production time windows.

The hybrid flow shop consists of the following stages:

- Stage 1: a single line including sequential processes such as, set-up of the line, cooking, and moulding the confectionary;
- Stage 2: the parallel drying cabinets. Stage 2 needs to consider the moulding duration as well. Stage 2 starts when the moulding of the intermediate starts and ends when the intermediate finishes drying.

The moulding time is taken into account in Stage 1 and in Stage 2. This does not mean we consider the moulding time twice. However, Stage 1 and Stage 2 do overlap during the moulding time as both the single line and one of the parallel machines need to be available during moulding.

We present the indices, parameters, and variables in Section 4.2.1. Section 4.2.2 discusses the constraints and Section 4.2.3 presents the objective we select.

4.2.1 Indices, Parameters and Variables

In this section we describe all the indices, parameters, and variables we use in our mathematical model.

Indices

$j, i \in J$ Jobs, including 2 dummy jobs, dummy job 1 at the beginning and dummy job 2 at the end of each schedule

$c \in C$ Drying cabinets

$p \in P$ Processing days

Parameters

m_j Moulding time of job j in hours

$v_{j,c}$ Drying time of job j in drying cabinet c in hours

$s_{i,j}$ Set-up time of job j after job i for Stage 1 in hours

k_j Cooking time for job j in hours

o_p First possible non-negative start time of processing day p (day 1 has as start time 0)

l_p Latest possible non-negative end time in hours of processing day p

$g_{j,c}$ Binary parameter, indicating whether job j, can dry in drying cabinet c ($g_{j,c} = 1$), or not ($g_{j,c} = 0$)

M Very large number

Variables

Non-binary variables

B_j Begin time of job j

T_j Duration of job j in Stage 1 (set-up, cooking, and moulding)

N_j Start moulding time of job j

E_j End time of job j in Stage 1, after moulding

$D_{j,c}$ End time of job j in cabinet c at Stage 2, after drying

Binary variables

$W_{j,p}$ Indicates whether job j starts processing at day p for $W_{j,p} = 1$, otherwise 0

$X_{i,j}$ Indicates whether job j is sequenced somewhere after job i in Stage 1 and therefore in Stage 2 as well for $X_{i,j} = 1$, otherwise 0

$Y_{j,c}$ Indicates whether job j is assigned to cabinet c for $Y_{j,c} = 1$, otherwise 0

$Z_{i,j}$ Indicates whether job j is sequenced directly after job i in Stage 1 for $Z_{i,j} = 1$, otherwise 0

Figure 4.1 presents the parameters (durations) and variables (time points) for Stage 1.

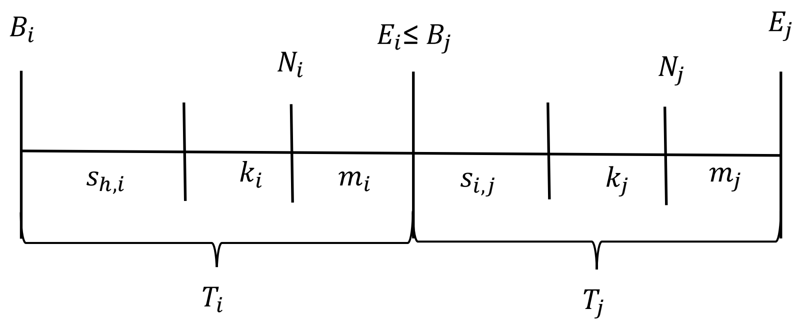

Figure 4.1: Two intermediates scheduled at Stage 1

In Figure 4.1 we schedule 2 jobs i and j in Stage 1. We schedule job i after a job h. Job i starts at time B_i and at this time the set-up from job h to job i starts with duration $s_{h,i}$. After the set-up, Brynild can start cooking job i that has a duration of k_i. Directly after cooking, the time point N_i indicates the start of the moulding process at which a drying cabinet must be available as well. The moulding has a duration of m_i. Directly after moulding at time point E_i, job i has completely entered the drying cabinet and job i exits the single line (Stage 1). The total time job i is in Stage 1 is T_i. At time point E_i, job j can start processing at Stage 1. The begin time of job j, B_j, must therefore start at the same time or later than the end time of the previous job, E_i. When job j starts, it goes through the same cycle as job i.

For job i, Stage 2 starts at timepoint N_i, see Figure 4.1. At time N_i one of the five drying cabinets needs to be available for job i. During the moulding time, the drying cabinet is emptied and filled again. Job i dries in a drying cabinet that is able to process this job. The drying time depends on the drying cabinet the job is assigned to.

4.2.2 Constraints

We summarize all the constraints in 20 mathematical (in)equalities. In this section we start with the constraints for Stage 1, then the constraints for Stage 2, and end with the general constraints.

Stage 1 constraints

Constraint [4.1.1] states that job j and job i process in series. We schedule job i before job j, or job j before job i, as only 1 of these statements can be true. Note that we cannot schedule a job before or after itself thanks to Constraint [4.1.2].

$$X_{i,j} + X_{j,i} = 1 \qquad \forall i,j \quad i \neq j \qquad [4.1.1]$$

$$X_{j,j} = 0 \qquad \forall i = j \qquad [4.1.2]$$

Constraint [4.1.3] makes sure that we never schedule job j directly after job i if we do not schedule job j somewhere after job i in general. These constraints also enforce 1 schedule, as we do not allow multiple sub-schedules. Without this constraint it could occur that we schedule job i directly after job j and schedule job j in general after job i. We want to enforce that $Z_{i,j}$ and $X_{i,j}$ have the same sequence.

$$Z_{i,j} \leq X_{i,j} \qquad \forall \text{i}, j \quad i \neq j \qquad [4.1.3]$$

We introduce 2 dummy jobs, which we locate at the start and the end of the production schedule. We need these dummy jobs to obtain values for $Z_{i,j}$. Without the dummy jobs we do not know, which job has no predecessor and, which job has no successor before we start scheduling.

Constraint [4.1.4] assures dummy job 1 is the first job in the schedule as we cannot schedule any other job before dummy job 1. Constraint [4.1.5] assures that dummy job 2 is last job as we cannot schedule any other job behind dummy job 2.

$$X_{i,j} = 0 \qquad j = dummy\,1 \quad \forall i \qquad [4.1.4]$$

$$X_{i,j} = 0 \qquad i = dummy\,2 \quad \forall j \qquad [4.1.5]$$

Constraint [4.1.6] ensures that we schedule each job directly after only 1 other job. The constraints leave out dummy job 1 for j as we schedule this job first. Constraint [4.1.7] ensures that we schedule each job directly before only 1 other job, except for dummy job 2.

$$\sum_{i \in J} Z_{i,j} = 1 \qquad \forall j \neq dummy\ 1 \quad i \neq j \qquad [4.1.6]$$

$$\sum_{j \in J} Z_{i,j} = 1 \qquad \forall i \neq dummy\ 2 \quad i \neq j \qquad [4.1.7]$$

The cooking time, moulding time, and the set-up time of job j together determine the duration of job j in Stage 1, see Constraint [4.1.8]. The set-up time is the summation over all jobs i that we can schedule before job j. We schedule only one of these jobs is directly before job j, as $Z_{i,j}$ has the value of 1 for only one specific job i. We add this set-up time to the duration of job j in Stage 1. Dummy jobs can induce set-up times at the beginning or end of the week when we desire.

$$T_j = k_j + m_j + \sum_{i \in J}(s_{i,j} * Z_{i,j}) \qquad \forall j \qquad [4.1.8]$$

Constraint [4.1.9] states that we calculate the end time of job j in Stage 1 by adding the duration of job j in Stage 1 to the begin time of job j.

$$E_j = B_j + T_j \qquad \forall j \qquad [4.1.9]$$

Constraint [4.1.10] ensures that jobs do not overlap in Stage 1. $Z_{i,j}$ is 1 if job j is the immediate successor of job i. In this case, the difference between the Stage 1 end time of job j and the Stage 1 end time of job i should be larger than the duration of job j at Stage 1. If we do not schedule job j directly after job i, $Z_{i,j}$ is 0, and this relation does not have to hold. These constraints apply to all jobs i and jobs j, however job i cannot be the same job as job j.

$$E_j - E_i + M(1 - Z_{i,j}) \geq T_j \qquad \forall i,j \quad i \neq j \qquad [4.1.10]$$

Constraint [4.1.11], [4.1.12], and [4.1.13] make sure that the jobs we schedule at Stage 1 stay within the production time windows. These constraints assure that every job j needs to start, and end within the same processing day. Constraint [4.1.11] states that we can schedule job j in only 1 time window.

$$\sum_{p \in P} W_{j,p} = 1 \qquad \forall j \qquad [4.1.11]$$

Constraint [4.1.12] ensures that when we schedule job j on day p, thus when $W_{j,p}$ is 1, the job does not start before processing day p starts. Because, when we schedule job j on day p, the begin time of job j, B_j, needs to be larger than or equal to the begin time of day p, o_p. Constraint [4.1.13] ensures that job j finishes before the end, E_j, of day p, l_p.

$$\sum_{p \in P}(W_{j,p} * o_p) \leq B_j \qquad \forall j \qquad [4.1.12]$$

$$E_j \leq \sum_{p \in P}(W_{j,p} * l_p) \qquad \forall j \qquad [4.1.13]$$

Stage 2 constraints

Constraint [4.1.14] assures that the start moulding time of job j in Stage 1 is the same or larger than the end time of job i in Stage 2. This is only the case when we schedule job i somewhere before job j, and we assign jobs i and j to the same drying cabinet. So, if we do not schedule job i somewhere before job j or if we do not assign job i and job j to the same drying cabinet, BigM ensures that the constraint is always fulfilled, independent of the values for $D_{i,c}$ and N_j.

$$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c}) \qquad \forall \mathrm{i}, j, c \quad i \neq j \qquad [4.1.14]$$

We indicate the moulding start time of job j by the end time of job j of which we subtract the moulding time, Constraint [4.1.15].

$$N_j = E_j - m_j \qquad \forall j \qquad [4.1.15]$$

We calculate the end time of drying job j in cabinet c, $D_{j,c}$, by adding the drying time to the end of job j at Stage 1. We only add the drying time when we assign job j to drying cabinet c. Otherwise, the end-drying time of job c is the same as the end time of job j in Stage 1, Constraint [4.1.16].

$$D_{j,c} = E_j + v_{j,c} * Y_{j,c} \qquad \forall j, c \qquad [4.1.16]$$

The fact that jobs get an end-drying time for cabinets to which we do not assign them, does not lead to further problems thanks to constraint [4.1.14]. We disregard these wrongfully assigned end-drying times with constraint [4.1.14] when we do not assign job j to drying cabinet c. Therefore, the end-drying time has no further influence on the other jobs.

Constraint [4.1.17] forces us to assign every job j to exactly 1 drying cabinet c, except for the 2 dummy jobs.

$$\sum_c Y_{j,c} = 1 \qquad \forall j \neq dummy\ 1\ and\ dummy\ 2 \qquad [4.1.17]$$

Constraint [4.1.18] ensures that we cannot assign job j to cabinet c, when the drying cabinet cannot process this job.

$$Y_{j,c} \leq g_{j,c} \qquad \forall j, c \qquad [4.1.18]$$

General constraints

Constraints [4.1.19] indicate the binary variables to be 0 or 1.

$$W_{j,p}, \quad X_{i,j}, \quad Y_{j,c}, \quad Z_{i,j} \qquad \in \qquad \{0,1\} \qquad [4.1.19]$$

Constraints [4.1.20] indicate that the different variables cannot become a negative value.

$$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c} \qquad \geq \qquad 0 \qquad [4.1.20]$$

4.2.3 Objective

Given these constraints, we can optimize the schedule of Brynild according to various objectives. Brynild proposes an objective to minimize the individual throughput time of each job to be able to produce more products per week. However, we do not select this objective, since minimizing the individual throughput time results in undesired effects on the schedule. Some undesired effects are drying as much as possible in the new drying cabinets and not necessarily starting Monday morning. We do not desire these 2 effects as Brynild regularly deals with unexpected changes, such as unplanned maintenance and additional production orders (that can maybe only dry in the new drying cabinets). Therefore, we desire to finish as early as possible in the week, so that Brynild can easier adapt the schedule through the week. When we schedule all orders at the end of the week we lose a lot of flexibility halfway through the week next to that, the probability increases that the schedule induces restrictions for the week after.

The objective we select for our model is to minimize the make span, see [4.1.21].

$$\min\,[\max_{j,c} D_{j,c}] \qquad [4.1.21]$$

We minimize the latest time that the last job finishes drying in one of the drying cabinets. The MILP enables us to schedule more quantity within a week than Brynild does now, which is Brynild's initial desire. This objective can deal best with the unexpected changes and limits the restriction for the upcoming week by finishing drying as early as possible. Our objective spreads the work force as evenly as possible over the drying cabinets to finish all jobs as early as possible. This objective makes it easier to observe whether Brynild can increase the throughput. In Appendix C we present a detailed description regarding the selection of the objective and the advantages and disadvantages of various objectives.

In Chapter 5 we introduce another objective solely optimizing Stage 1. We still consider Stage 2, however we do not optimize Stage 2. We use this model to quantify the effect of different objectives.

4.3 HEURISTICS BASED ON OUR DEVELOPED MODEL

In the MILP from Section 4.2 we consider multiple decision variables, resulting in a large solution space. Therefore, we expect (and observe during the experiments from Chapter 5) that the computation time of the MILP is relatively high.

We base the heuristics on the MILP and not on (meta-)heuristics, such as local search, tabu search, or simulated annealing. We prefer to conduct further research regarding the possibilities of the MILP and how heuristics based on this MILP perform rather than developing a (new) heuristic, as there is very limit literature research based on this problem. We cannot apply a (meta-)heuristic without fully adapting the heuristic to our problem. During the adaption, the probability is high that (unknown) obstacles emerge. For example, when an operator for the (meta-)heuristic does not consider infeasible solutions, then escaping from a local optimum becomes very difficult. It is difficult to escape a local optimum as a local search (meta-)heuristic has to 'navigate' through large spaces of infeasible solutions to obtain a good solution. When the drying cabinet for a job on Stage 1 is not available the job jams the whole production line until the drying cabinet is available. When the job on Stage 1 is jamming the schedule too long we cannot schedule all jobs within the week, and the schedule becomes infeasible. We could temporarily discard the end time of the week to create more feasible schedules, which makes it for the operator of the heuristic easier to navigate. If we allow to temporarily discard the end time, we need to 'remember' much more schedules to later determine the best feasible schedule including the weeks end time. Next, we mislead the operator partially, therefore the operator uses more computation time to obtain schedules that we later find infeasible. This approach results in a larger solution space. To thoroughly navigate this larger solution space that includes a lot of infeasible schedules the probability can arise that these heuristics require more time and memory than solving the MILP. Therefore, we do not focus on these (meta-)heuristics in this thesis.

We find it useful to observe whether the MILP is still applicable as a basis when the problems require faster solving or if the problems become larger. Therefore, we introduce these 3 heuristics we base on our main model from Section 4.2:

- Our MILP without sequence dependent set-up times;
- Our MILP that we decompose in an assignment and a sequencing MILP, which we compute sequentially;
- Our MILP with a maximum computation time of 10 minutes.

We expect that these heuristics achieve (near) optimal solutions in less time than the main model. In Section 4.3.1 we describe the MILP without SDST, in Section 4.3.2 the decomposition of our MILP, and in Section 4.3.3 the heuristic where we reduce the computation time.

4.3.1 MILP without SDST

Brynild indicates that most of the time they do not consider the sequences dependent set-up times while determining the week schedules, and assume for scheduling purposes a set-up time of 2 hours, $s = 2$. Therefore, we decide to change the variable time of job j at Stage 1, which includes the SDST, to a parameter including the standard set-up time of 2 hours.

We change the duration at Stage 1 for each job from a variable to a parameter therefore, we no longer need the dummy jobs to describe the MILP. Next, we discard the variable $Z_{i,j}$ that we use to determine whether we schedule job j directly after job i. Therefore, we need to change and discard multiple constraints and variables from our main MILP. Thus, we make the following 3 changes:

1. We no longer need the dummy jobs therefore, we remove all constraints regarding the dummy jobs from model, which are [4.1.4] and [4.1.5];

2. We no longer need variable $Z_{i,j}$ therefore, constraint [4.1.19] can discard the variable $Z_{i,j}$. Next, we remove the following constraints including $Z_{i,j}$ from the model [4.1.3], [4.1.6], and [4.1.7]. Constraint [4.1.10], includes $Z_{i,j}$ as well and for this constraint the variable $Z_{i,j}$ changes in $X_{i,j}$;

3. Variable T_j, the duration of job j in Stage 1, becomes a parameter we denote as t_j. So, equation [4.1.8] changes into equation [4.2.1].

$$t_j = k_j + m_j + s \qquad \forall j \qquad [4.2.1]$$

We present the adaptions to the MILP as Model 7 in Appendix D. We base the order of Appendix D on Table 5.8 from Chapter 5.

4.3.2 Two separate MILPs

The second heuristic is decoupling the assignment of jobs to drying cabinets and the sequencing of jobs in Stage 1. Similarly to the heuristic from Section 4.3.1, we base this heuristic on turning a variable in a parameter to reduce the solution space. In the initial model we conduct the assignment and sequencing of jobs simultaneously, which is interactive. In this heuristic we use 2 separate MILPs. The first MILP assigns the jobs to cabinets, like GAP in Section 3.2.1. As objective we decide to spread the workload as evenly as possible over the drying cabinets by minimizing the maximum total time that we assign to the drying cabinets, see objective [4.3.1].

$$\min\,[\max_{c} R_c] \qquad [4.3.1]$$

Where R_c is the total time in quantity we assign to cabinet c. The total time is the summation of the drying time and the moulding time, see equation [4.3.2].

$$R_c = \sum_j (Y_{j,c} * (m_j + v_{j,c})) \qquad \forall c \qquad [4.3.2]$$

We select this objective as we do not know exactly how full a drying cabinet can be. We present the decision for our objective of this assignment model in part 1 of Model 9 in Appendix D. After determining the assignment of jobs to drying cabinets we use this output as input for our MILP to decide the sequence in Stage 1. The Variable $Y_{j,c}$ indicating whether we assign job j to drying cabinet c becomes the parameter $y_{j,c}$, using the output of the assignment MILP. Therefore, we no longer need constraints [4.1.17] and [4.1.18] in the sequencing MILP as we already use these constraints in the assignment MILP. We present the sequence MILP in part 2 of Model 9 in Appendix D.

4.3.3 10 min MILP

In the third heuristic the MILP of Section 4.2 has a maximum computation time of 10 minutes. We do not base this third heuristic on reducing variables and constraints, as opposed to the two previous heuristics we describe. We base this heuristic on the observation that our MILP model typically finds a good, or even optimal, solution within the first few minutes. The additional time the solver uses is to eliminate the remainder of possible solutions to state with certainty that the already found solution is optimal. Brynild indicates that they like to use the model as a scenario tester as well. Therefore, the computation time should not be more than 1 hour. However, we observe in Chapter 5 that 10 minutes is for most instances already enough to create (near) optimal schedules.

We present the MILP from Section 4.2 without explanation as Model 1 in Appendix D.

4.4 CONCLUSION

Our mathematical model focuses on the Støperi 1 line as this line contains the drying area, which is the bottleneck. We do not take the oiling and sanding into account as this process has enough production capacity to always process the output from demoulding the intermediates, therefore oiling and sanding do not become a bottleneck when the bottleneck shifts. Given this focus, we model the scheduling problem as a two-stage hybrid flow shop in which the drying cabinets are non-identical parallel machines. The model includes sequence dependent set-up times, production time windows, and no buffer capacity between the 2 stages.

We assume that we structure input data in such a way that jobs consist of 1 or more intermediates, which can dry together and, in terms of quantity always fit in 1 drying cabinet. Another assumption we make is that the drying cabinets are available at the beginning of the week. In our mathematical model we focus on the end time of the drying cabinets and not on the time we empty the drying cabinets. The other additional assumptions we make are:

- All machines in the production line are available at all times; no breakdowns or maintenance;
- Jobs process always without error;
- We cannot interrupt job processing, so no pre-emption;
- We cannot process jobs on more than 1 machine simultaneously;
- A machine may only process 1 job at a time;
- All data is known deterministically when creating the schedule;
- We do not consider set-up time between weeks.

In addition to our main MILP model with as objective to minimize the make span of the schedule, we introduce 3 heuristics we base on this main MILP:

- Our MILP without sequence dependent set-up times;
- Our MILP that we decompose in an assignment and a sequencing MILP, which we compute sequentially;
- Our MILP with a maximum computation time of 10 minutes.

We do not research meta-heuristics further as Brynild has a 2-stage HFS problem with no buffer capacity between the 2 stages and production time windows. These two conditions make it difficult to apply and adapt meta-heuristics to solve the problem. Furthermore, we find it useful to observe whether the MILP is still applicable as a base when the problem requires less computation time or becomes larger.

In Chapter 5 we evaluate the experiments regarding the MILP we develop in Section 4.2, the 3 heuristics, and other variants of our MILP.

5 EXPERIMENT DESIGN & RESULTS

In this chapter we answer the research question: '*How does the proposed scheduling method perform?*'. In Section 5.1, we compare the two-stage hybrid flow shop with non-identical parallel machines, sequence dependent set-up times, production time windows, and zero buffer points MILP to Brynild's week schedules. Thereafter, in Section 5.2 we compare the test results of the 3 heuristics from Section 4.3 to the results from Section 5.1. In Section 5.3 we examine the general capabilities of our MILP and various other MILPs that we base on zero buffer capacity between 2 stages.

We compute all instances within this research with the CPLEX 12.10 solver. We conduct the experiments with a Dell Precision M2800, with a x64-based processor of 2.50GHz and a RAM of 8.00 GB.

5.1 CASE STUDY: BRYNILD

Section 5.1.1 describes the experiment design. Section 5.1.2 summarizes and evaluates the results of the experiment design.

5.1.1 Case study: Experiment Design

In the Brynild case study, our MILP model reschedules Brynild's Week 2 to Week 8 from 2020. We present the input data of these weeks in Appendix E. We perceive the weeks as independent, due to the buffer function of the weekends, as we explain in Section 4.1.3.

Since we reschedule independent weeks, we make use of a 1-week scheduling horizon. We use a 1-week scheduling horizon for two main reasons. Firstly, this 1-week horizon represents current practices at Brynild best, as Brynild currently schedules one week at a time as well. Secondly, the output of the planning indicates the intermediates and their quantity that Brynild needs to produce in a 1-week period. If we consider scheduling more than one week at the time, the solution space becomes larger and when creating the best schedule the probability is high that intermediates change between weeks. However, by interchanging the intermediates between weeks, we nullify the calculations of the planning regarding holding costs, operating costs, etc., since we do not produce the correct intermediates in the correct week any longer.

We start with the determination of jobs in Section 5.1.1.1. In Section 5.1.1.2 we discuss how we obtain the input data for the jobs. In Section 5.1.1.3 we explain the experiment design.

5.1.1.1 Determination of the jobs

The raw data available from the planning process is not, like we assume in Chapter 4, structured in such a way that jobs consist of 1 or more intermediates, which can dry together, and that the jobs fit in 1 drying cabinet. This assumption requires data pre-processing. We schematically present the pre-processing in the flowchart in Figure 5.1.

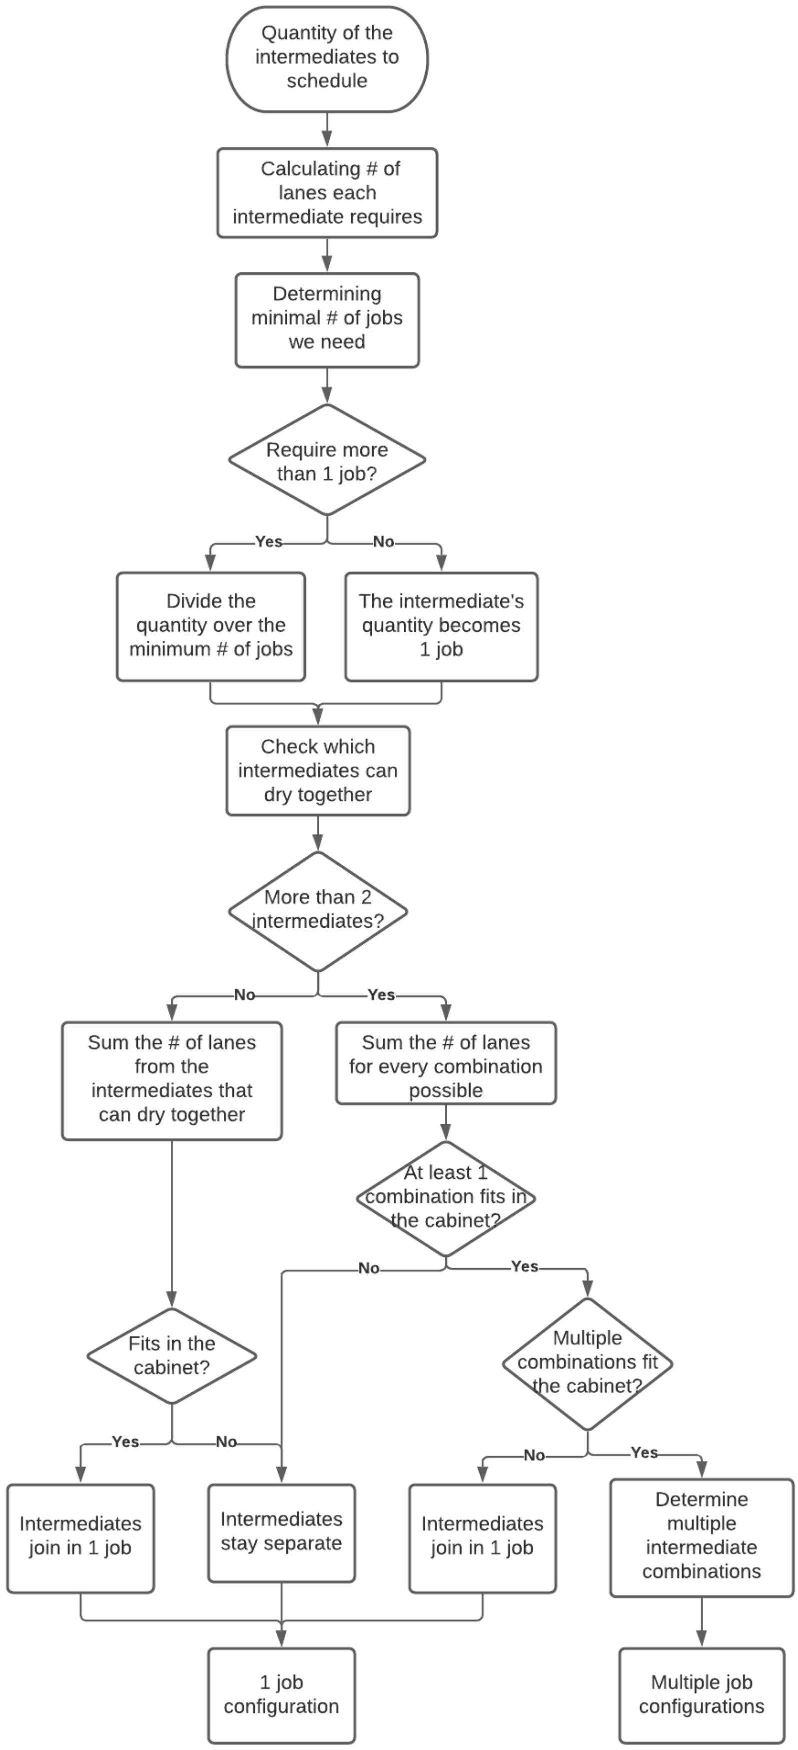

Figure 5.1: Flowchart of the pre-processing of the input data

The flowchart in Figure 5.1 starts with the quantity of each intermediate that Brynild plans to produce in one week. We calculate the number of lanes that each intermediate requires to dry the quantity. With this information we calculate the minimum number of jobs we need for 1 intermediate to dry all the quantity. We calculate the minimum number of jobs we need by rounding up: dividing the number of lanes for an intermediate by the number of lanes available per drying cabinet. When we require more than one job to produce the quantity of the intermediate, we divide the quantity over the minimum number of jobs. How we divide the quantity over the various jobs, is up to the decision maker. We propose 3 different heuristics that the decision maker can use to divide the quantity:

1. Divide the quantity evenly over the minimum number of jobs;
2. Fill the jobs completely and the last job contains the rest quantity;
3. Obtain a variety of different job sizes when we require more than 2 and less than 6 jobs, as we do not encounter an intermediate that requires more than 5 jobs. We fill the first job completely, the second job consists of the quantity that a job would have when there is an even distribution, for the third job we subtract the difference between the first and second job from the quantity of the second job, and the fourth job is the average between the first and second job. When there is not enough quantity left to construct a job, the job just consists of the rest quantity.

Figure 5.2, illustrates the 3 constructive heuristics.

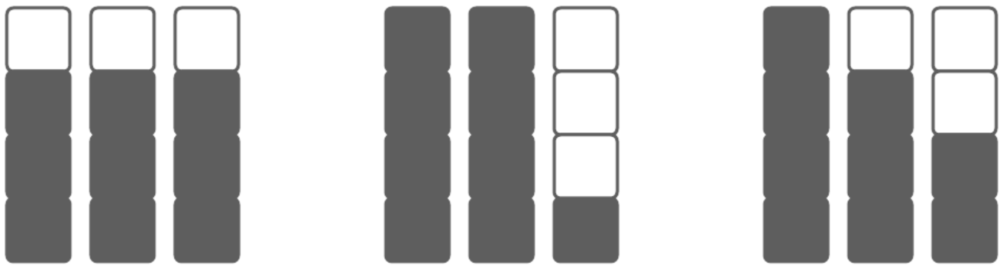

Figure 5.2: Heuristics for the quantity division over jobs

Also, we could use other heuristics, which generate more than the minimum number of jobs. For example, 'every job can only be half full'; in this case we need 5 jobs, if we consider the same situation as we illustrate in Figure 5.2. In the remainder of this research, we only consider the 3 heuristics we describe above and in Figure 5.2, since these lead to the minimum number of jobs. We make this decision because the model requires more computation time when we use more jobs as input.

After determining the minimum number of jobs we need for each intermediate, we check whether certain intermediates can dry together. When an intermediate can dry with more than one other intermediate, we determine for every combination of previously made jobs if they fit in the drying cabinets together. We combine the jobs consisting of intermediates that can dry together and do fit in the same drying cabinet as 1 job. When we combine intermediates in 1 job, we add the additional time of cooking, moulding and set-up between the intermediates to the job's moulding time. We do so because both the drying cabinet and the production line need to be available when producing the additional intermediates. When the intermediates that can dry together do not fit in the same drying cabinet, the intermediates remain separate jobs. It may occur that multiple intermediates can dry together, and multiple combinations of these intermediate batches fit together in 1 drying cabinet. In that case, the result of our data pre-processing is 2 or more job configurations, thus 2 or more input data sets. By creating multiple job configurations and optimizing each of them with our MILP, the overall solution space becomes larger and the probability of obtaining the optimal solution increases. The disadvantage is that the computation time increases as well.

We automate this process of determining jobs in Excel, except for the part where the answer to the question ‘multiple combinations fit the drying cabinet’ is ‘yes’. We do not program this part in Excel as this situation does not occur in our sample weeks and is not necessary to include in our experiments. The process itself requires a few seconds. Therefore, the disadvantage of pre-processing requiring more time is negligible.

5.1.1.2 Input data

We present the 7 weeks we use to conduct our experiments in Table 5.1.

Table 5.1: Brynild’s week descriptions

Week	# of jobs	# of machines	Original schedule
2	11	5	3-shift
3	13	5	3-shift
4	11	5	3-shift
5	11	5	3-shift
6	10	5	2-shift + Sat
7	9	5	2-shift + Sat
8	10	5	2-shift+ Sat

The number of jobs varies between 9 and 13. We consider 5 drying cabinets, which are all available in the weeks under consideration. Therefore, there is no variation in the number of machines in this case study. In Section 5.3 we do vary the number of parallel machines. Brynild originally schedules four of the seven weeks using a 3-shift schedule, and the other three with a 2-shift + Saturday schedule.

In Chapter 2, we explain that a 3-shift schedule starts production on Sunday at 22:30h and finishes on Friday at 14:30h. Therefore, in case of a 3-shift schedule there is 1 large time window for index p. In the schedule we create with our MILP, the first job has the start time 0, which indicates Sunday 22:30h. In a 2-shift schedule the first shift of the week starts on Monday at 6:30h. So, start time 0 in a 2-shift schedule indicates Monday 6:30h. Within the 2-shift schedules the production in Stage 1 stops during the night. For example, we do not allow the model to schedule any jobs between time l_p= 16 and o_p= 24, which indicates Monday 22:30h and Tuesday 6:30h. We present the specific values we use to obtain a 2-shift and a 3-shift schedule in Appendix E.

The set-up time depends on the tools that the intermediates on the production line require, as we describe in Chapter 2. When the intermediates make use of the same tool, there is a set-up time of 1.5 hours, and when the intermediates make use of different tools there is a set-up time of 2 hours. For jobs containing the same intermediate there is zero set-up time. The cooking time is 0.5 hours for every job. We present the drying time and tools for each of the intermediates in Appendix E.

We calculate the moulding time in minutes using the following equation:

$$Moulding\ time\ in\ minutes = \frac{\#Pallets * 150}{27}$$

There are 150 trays on 1 pallet, and on average it takes 1 minute to mould 27 trays. We require the moulding time in hours, thus we divide this value by 60 minutes to obtain the value for our moulding time parameter. In case 2 or more intermediates dry together we add the additional time for set-ups and cooking to the moulding time as well.

In Appendix E we state the intermediates and the quantity in kg that Brynild produces in weeks 2 to 8.

5.1.1.3 Experiment design

The experiments we conduct analyse 3 specific aspects: constructive heuristic for the input data selection, multi-objective evaluation, and shift evaluations. We start with testing, which constructive heuristic within the pre-processing is best to use for the further experiments. We conduct an experiment using our MILP and the 7 weeks from Brynild and test the different heuristics for every week. Therefore, we conduct 21 tests. From these tests we collect the objective value and the computation time of the model. We consider the computation time as well, as different job sizes can have different effects on the computation time, especially in a 2-shift schedule.

After determining, which constructive heuristic we use for further experiments we test 2 objectives. The first objective we already introduce in Chapter 4, which is the end-drying objective: the minimization of the make span (after drying):

$$\min [\max_{j,c} D_{j,c}],$$

where $D_{j,c}$ is the end-drying time of job j in cabinet c.

The second objective, the Stage 1 objective, is the minimization of the production make span of all jobs at Stage 1:

$$\min [\max_{j} E_j],$$

where E_j is the end time of producing job j in Stage 1.

We consider the second objective as only in Stage 1 Brynild requires staff to work the production line. When the jobs are inside the drying cabinets Brynild can deploy the personnel somewhere else in the production lines. When using the second objective we still consider the constraints for the drying cabinets, and the difficulties that come with scheduling the drying cabinets. However, we only optimize the make span of Stage 1 and not the entire HFS problem; we foresee that the model schedules the jobs with longer drying times last. The Stage 1 objective solely optimizes Stage 1, therefore we do not select this second objective as our main objective for our MILP. However, we do want to conduct experiments with the second objective since the results could be interesting for Brynild whenever they want to reduce personnel costs. Next, we aim to quantify the effect of the different objectives for Brynild.

We conduct the experiments using the 2 objectives for all 7 weeks. From these tests we collect the objective value, the value for the other objective, the set-up time, and the computation time of the model. We compare the objective values of both objectives to Brynild’s schedules. Thereafter, we evaluate the influence of the set-up times on the Stage 1 end times.

The final experiment we conduct is to analyse whether it is possible to ‘scale down’ the number shifts we use in the schedules. With scale down we mean scheduling all intermediates that Brynild originally schedules in a 3-shift schedule in a 2-shift + Saturday schedule or even in a 2-shift schedule. For the 2-shift + Saturday schedules we examine whether it is possible to schedule all the intermediates without the Saturday. Due to time constraints, we conduct these 18 experiments only for the drying objective as this is our main objective.

5.1.2 Case Study: Results MILP model

We start with evaluating the results of the constructive heuristics in Section 5.1.2.1. Thereafter, we evaluate the results of the experiments we conduct with the first objective in Section 5.1.2.2, and with the second objective in Section 5.1.2.3. In Section 5.1.2.4 we evaluate the influence of the set-up times on the Stage 1 end time. Finally, in Section 5.1.2.5 we assess the possibility of scaling down the original schedules of Brynild.

5.1.2.1 Constructive heuristics

In Section 5.1.1 we describe 3 heuristics for dividing the quantity of intermediates, which are too large for 1 job. Every week has 1 or more intermediate(s) of which the quantity requires more jobs. Week 4 and Week 5 have only intermediates that require a maximum of 2 jobs. Thus, for Week 4 and Week 5 heuristics 2 and 3 result in the same input instance, as both heuristics start filling one job completely and the second job obtains the rest. When the model requires more than 2 hours to complete the schedule, we interrupt the model. The reason for this interruption after 2 hours is that we have limited time available for this research. Our main focus is not to find the best constructive heuristic, but rather on developing a scheduling method after the pre-processing. Luckily, we only have to interrupt the model in Week 3.

Figure 5.3 presents the objective values we find using the 3 heuristics.

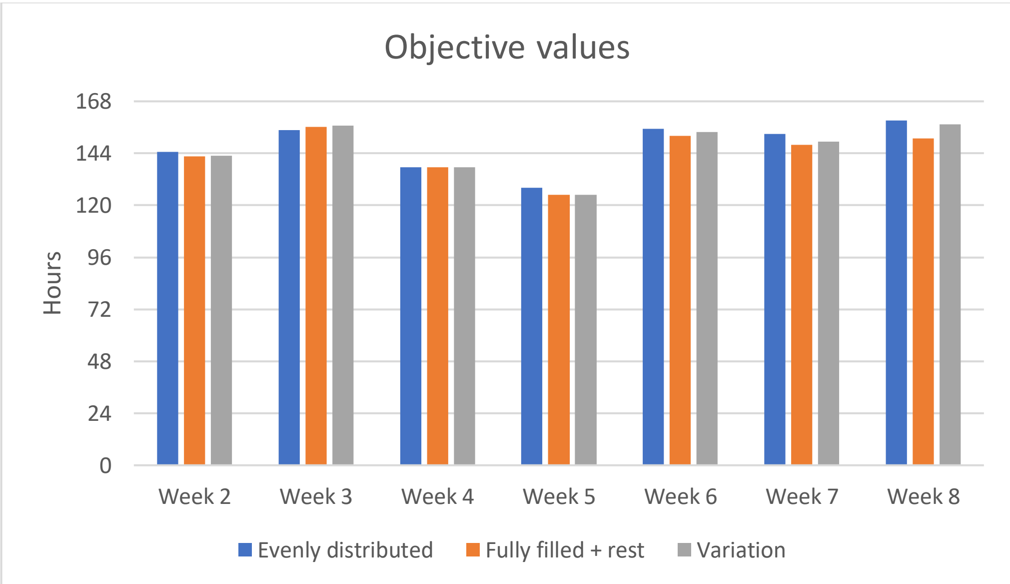

Figure 5.3: Objective values for the 3 constructive heuristics

For every week the second heuristic achieves the same or a better objective function than the third heuristic. In Week 3, the first heuristic achieves the best objective value. If we solely consider the objective values, the second heuristic would dominate the third heuristic. However, we consider the computation time as well. So, when the second heuristic computes the schedule faster than the third heuristic for every week, we dismiss the third heuristic completely, as the second heuristic then dominates the third heuristic in both aspects.

Figure 5.4 presents the time in seconds to find the schedule for the 3 heuristics.

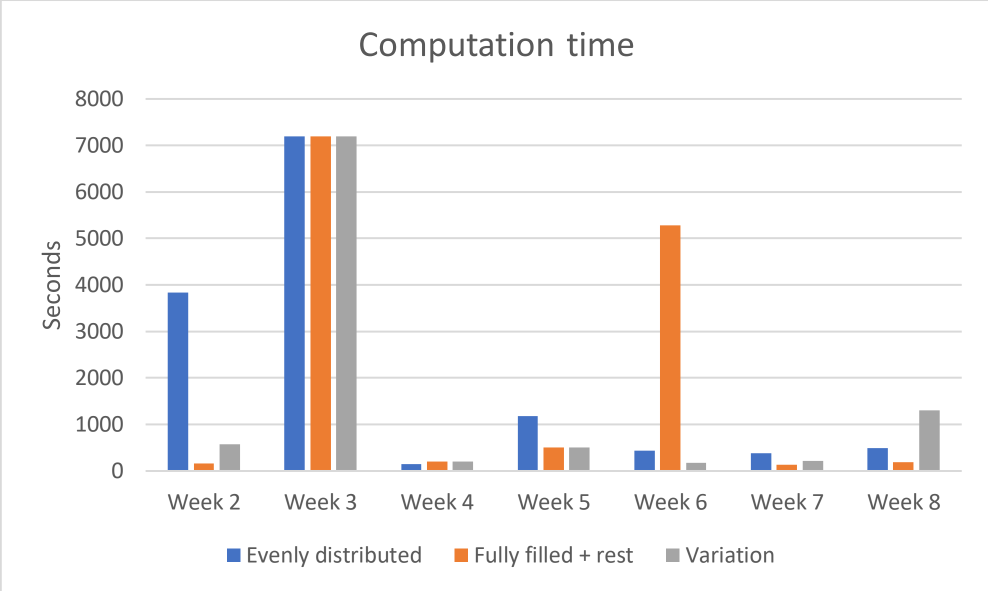

Figure 5.4: Time in seconds to find the objective

In Figure 5.4, we observe that in Week 6 the second heuristic requires the most time. Therefore, the second heuristic does not dominate the third rule and we cannot discard any of the heuristics. Week 3 is the only week that requires more than 2 hours to obtain the objective values. Each of the heuristics requires in at least one of the weeks the most computation time. The first heuristic requires the most computation time in weeks 2, 5, and 7, the second heuristic in Week 6, and the third heuristic in Week 8.

Our sample size consists of a limited period of 7 weeks. The experiments suggest that the second heuristic: 'Fill the jobs completely and the last job contains the rest quantity', has the most potential. We do not analyse the outliers within the computation time, since our main focus in on our MILP from Chapter 4, and not on the pre-processing of the jobs. We mainly use these experiments to indicate further research for Brynild.

We decide to conduct the rest of the experiments using the first constructive heuristic: 'Divide the quantity evenly over the minimum number of jobs'. This heuristic is most similar to the division of intermediates that Brynild currently uses. If the experiments show that 1 heuristic is superior in every week for both the objective value and the computation time we would use that heuristic. However, currently none of the heuristics performs superior in every week, therefore we consider a trade-off between the second heuristic with the highest potential and the first heuristic, which is the default heuristic that creates the most similar jobs when we compare them to Brynild's original job configuration. The second heuristic shows great prospect, however we do not want to contaminate possible scheduling output by deviating so much from Brynild's original job configuration in the pre-processing. If we had enough time, we would conduct all upcoming experiments with both job configurations to obtain more data however, for now we find the data too limited to choose a heuristic that deviates much from the original, without knowing the full extend and reasons behind the results from Figure 5.4. We recommend researching the outliers before switching the heuristic for dividing the quantity over jobs, as the reason behind the largest outlier in Week 6 can be structural in one way or another.

Next, working with the first heuristic emphasizes the effect of scheduling the jobs with our MILP and shows the potential of our method in comparison to Brynild's current scheduling practice, which, for now, we prefer over including the effects of the pre-processing.

5.1.2.2 *Min max end-drying time for the original schedules*

As we explain in Section 5.1.1.3, we conduct several experiments to compare multiple objectives. In this section we evaluate the end-drying objective. On average, we reduce the maximum end-drying time by 20.5%. Figure 5.5 presents the maximum end-drying time of Brynild's schedules and the schedules of our MILP model for each week.

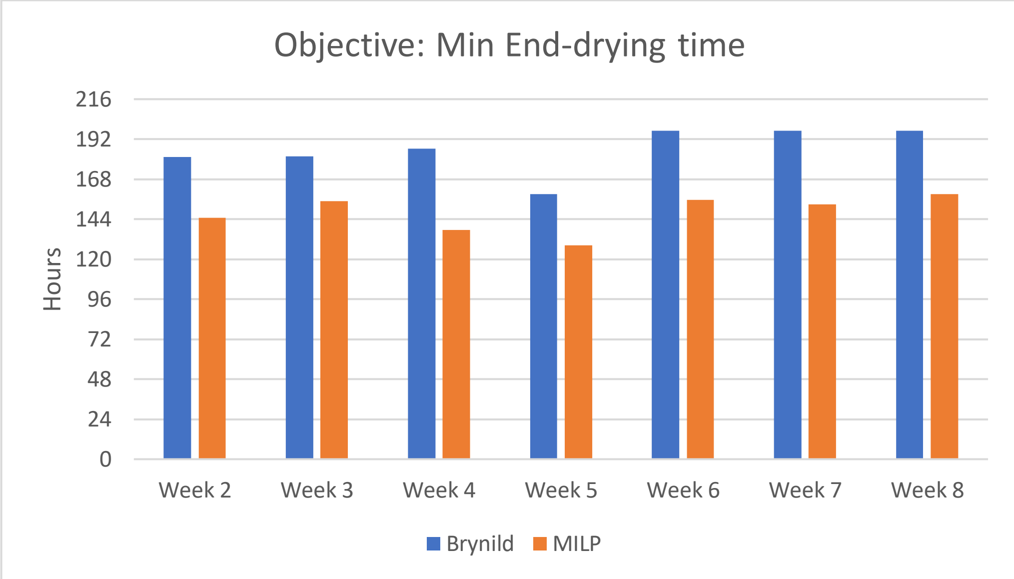

Figure 5.5: End-drying time comparison between Brynild and our MILP

From Figure 5.5, we conclude that our model gives a reduction in the latest end-drying time for each week. When we consider the reduction in hours, on average the latest drying cabinet finishes 38 hours earlier. For Week 3 we achieve the least reduction in hours, namely 27 hours, nevertheless, this is still more than 1 day. For Week 4 we achieve the largest reduction in hours, namely 49 hours, which is more than 2 days faster than Brynild's original schedule. The difference between Brynild's end-drying times and the end-drying times of our scheduling method is significant ($p<0.05$, see Appendix G).

Focusing on minimizing the end-drying times, has as a consequence that we consider the end time of Stage 1 only as a constraint; for five of the seven weeks the Stage 1 end time increases. The maximum increase of the Stage 1 end time is 5.87%, which is in this case approximately 7 hours.

Appendix F gives more information regarding the test results of our MILP, including the end time of Stage 1 we achieve under the end-drying objective in comparison to Brynild's end Stage 1 time.

5.1.2.3 Min max end time Stage 1 for the original schedules

In this section we evaluate the Stage 1 objective. On average, we reduce the maximum Stage 1 end time by 15.5%. Figure 5.6, presents the maximum Stage 1 end time of Brynild's schedules and the schedules of our MILP for each week.

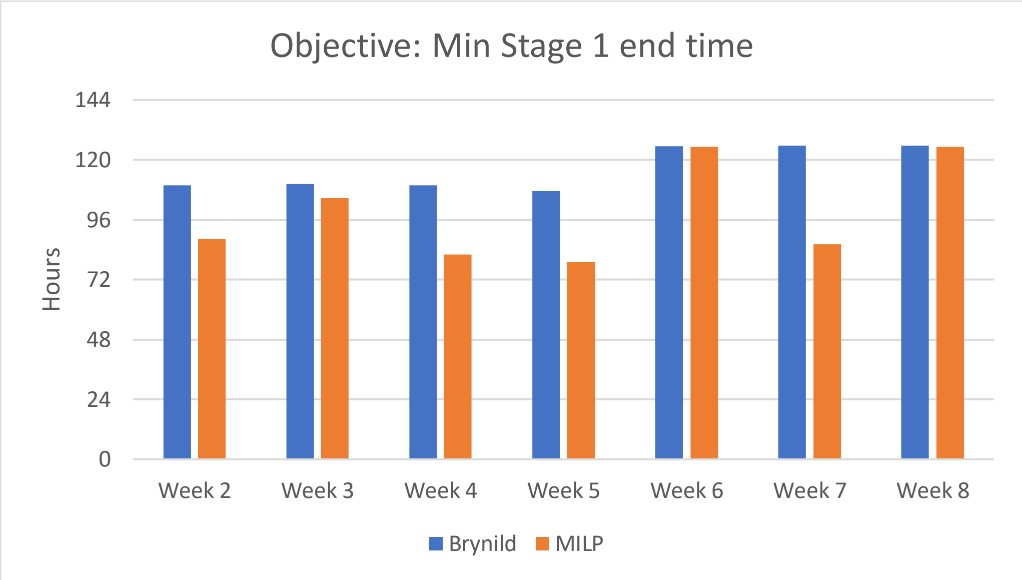

Figure 5.6: End Stage 1 time comparison between Brynild and our MILP

From Figure 5.6, we conclude that our model gives a reduction in latest Stage 1 end time for each week, including Week 6 and Week 8 that we schedule using a 2-shift + Saturday schedule. The objective values of weeks 6 and 8 do not differ much from Brynild's original schedule. In Week 7 however, Stage 1 finishes far earlier than in Brynild's schedule. This shows that we probably do not need the Saturday to schedule all intermediates. We discuss this further in Section 5.1.2.5, where we consider the reduction in shifts. On average the last job of Stage 1 finishes 17.5 hours earlier than in Brynild's schedules. The schedule of Week 6 shows the least reduction in hours, namely 0.3 hours. The schedule of Week 7 has the largest reduction in hours, namely 39 hours, which is a reduction of more than 1.5 days in comparison to Brynild's schedule.

Focussing on the Stage 1 end time, we disregard the drying end times. For six of the seven weeks the drying end time decreases as well; with a maximum of 19.84%. However, for three of the seven weeks the last drying cabinets do not finish before the next week starts. In Week 6 the last drying cabinet even finishes on Tuesday at 23:38h the following week.

Appendix F gives more information regarding the test results of our MILP, including the end-drying time we achieve under with the end Stage 1 objective in comparison to Brynild's end-drying time.

5.1.2.4 *Evaluation set-up times for each objective*

After evaluating the 2 objective values in the previous sections, we continue with analysing the set-up times we achieve for both objectives and present them in Figure 5.7.

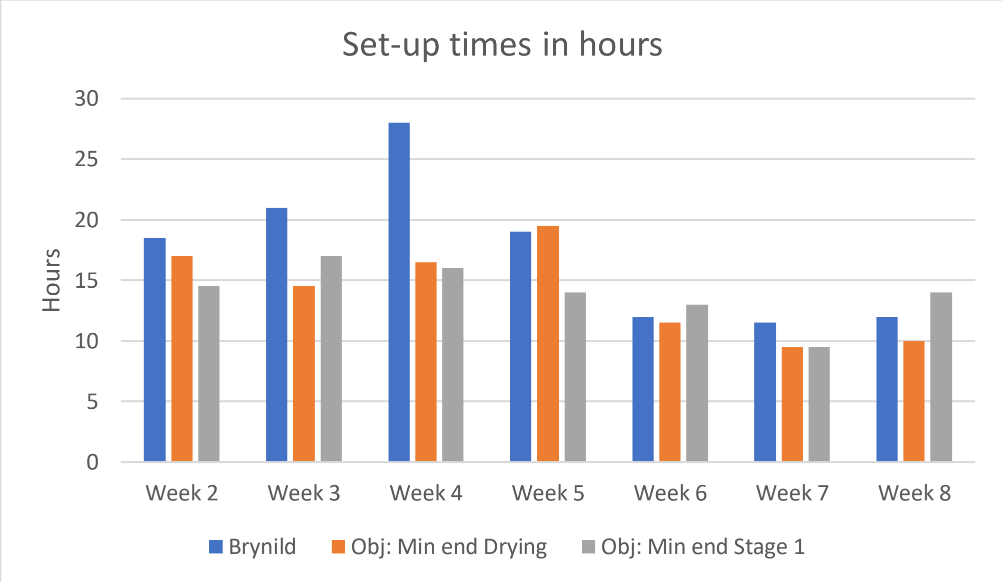

Figure 5.7: Set-up times of Brynild and both objectives for each week

Figure 5.8 presents the end time of Stage 1 for each week and objective.

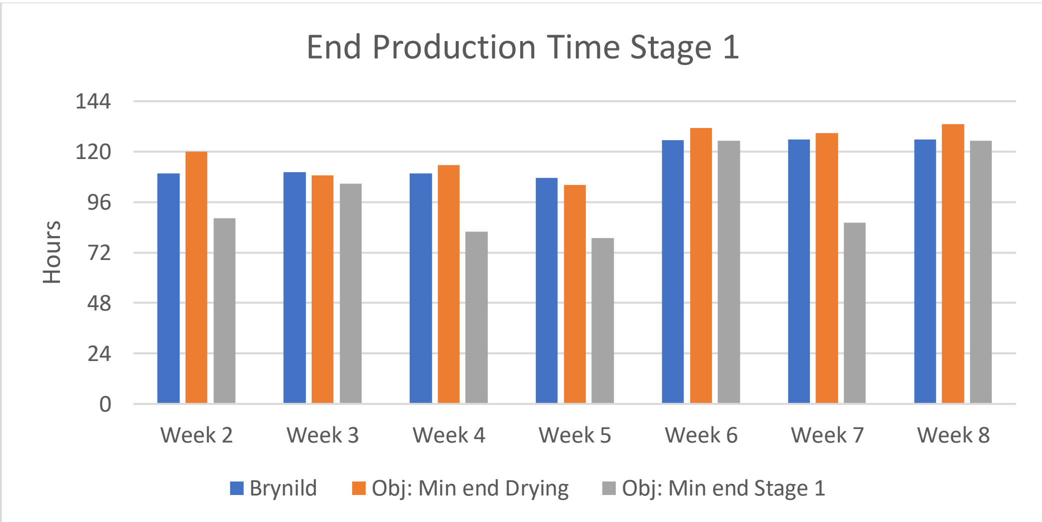

Figure 5.8: End production times of Stage 1 of Brynild and both objectives for each week

The sequence dependent set-up time is the only variable in the processing time of a job at Stage 1. Therefore, more set-up time leads to more total production time at Stage 1. Thus, it is plausible to assume that more total production time, probably leads to a later end time of Stage 1. After analysing both Figure 5.7 and Figure 5.8 we conclude that more set-up time does not necessarily mean a later end time at Stage 1. We observe that in none of the weeks the expectation, that more set-up time means a later Stage 1 end time, upholds. For example, in Week 6, the objective to minimize the end time of Stage 1 (in grey) has the most set-up time however, finishes the earliest. The same applies for Week 8. We conclude that less sequence dependent set-up time does not necessarily come with an earlier finish time of Stage 1. As there is no other variable in the Stage 1 production time, we hypothesise that the availability of drying cabinets in Stage 2 has a larger effect on the end time of Stage 1 than the set-up times.

5.1.2.5 *Scaling down the original schedule sizes*

For the schedules of weeks 2 to 5 we try to reduce the 3-shift schedule to a 2-shift + Saturday schedule and if possible, to a 2-shift schedule. For the schedules of weeks 6 to 8 we try to reduce the 2-shift + Saturday schedule to a 2-shift schedule. We present our results in Table 5.2.

Table 5.2: End-drying time for the original schedules and the reduced schedules

Week	Original Brynild schedule	Reduced schedule	Original end-drying time	New end-drying time (old schedule)	New finished drying (reduced schedule)
2	3-shift	2-shift + Sat	One week later Monday 11:43h	Saturday 23:15h	Sunday 1:50h
3	3-shift	3-shift	One week later Monday 12:18h	Sunday 9:11h	
4	3-shift	2-shift + Sat	One week later Monday 16:49h	Saturday 15.58h	Sunday 17:27h
5	3-shift	2-shift + Sat	Sunday 13:43h	Saturday 6:44h	Sunday 11:50h
6	2-shift + Sat	2-shift + Sat	One week later Tuesday 11:25h	Sunday 17:53h	
7	2-shift + Sat	2-shift	One week later Tuesday 11:43h	Sunday 15:20h	Sunday 15:20h
8	2-shift+ Sat	2-shift + Sat	One week later Tuesday 11:43h	Sunday 21:41h	

The third column 'Reduced schedule' in Table 5.2 presents the schedule to which we reduce the original schedule from column two. We observe that we can schedule the intermediates of weeks 2, 4, 5, and 7 in a schedule with less processing time for Stage 1. The fourth column 'Original end-drying time' presents on what day and time Brynild's original schedule finishes drying. All weeks, except for Week 5, cause a restriction for the next week as 1 or more drying cabinets do not finish drying before the start of the upcoming week. The column 'New end-drying time (old schedule)' presents the latest end-drying time of our MILP using the original number of shifts. All drying cabinets finish before the start of the upcoming week, and do not generate any restrictions for the next week. This is in line with our assumption of independent scheduling weeks from Chapter 4. When we use the reduced shift-schedules and less time for Stage 1 is available, the end-drying time of all weeks still finishes before the start of the upcoming week. Thus, we conclude that these reduced schedules do not generate any disadvantages, in terms of constraints for the upcoming week. On top of that, the reduced schedules require less time in Stage 1, and therefore lower personnel costs in these 4 weeks. In Appendix G we examine why we can or cannot scale down certain weeks. As a result, we expect that scaling down from a 3-shift schedule to a 2-shift + Saturday schedule is possible when the average production hours at Stage 1 is under 80 hours, which equals the number of hours available at Stage 1 in a 2-shift + Saturday schedule. Next, we expect that scaling down from a 2-shift + Saturday schedule to a 2-shift schedule is influenced by the available drying time and the total drying time required by the jobs, rather than by the available time at Stage 1.

We give a graphical representation of the schedule of Week 5 using a 3-shift schedule: in Figure 5.9 we present Stage 1 and in Figure 5.10 we present Stage 2.

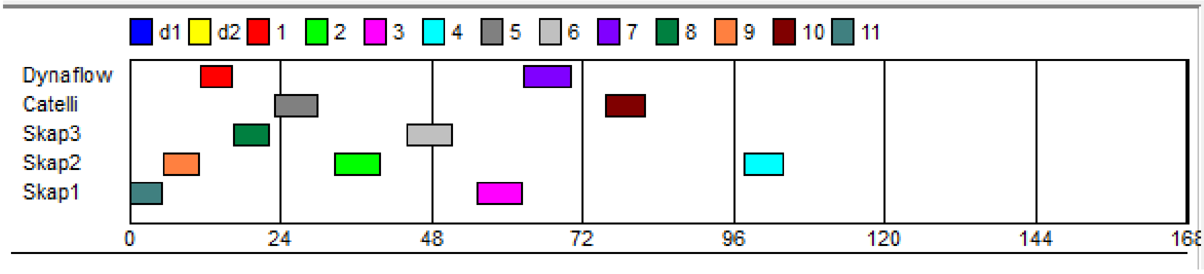

Figure 5.9: A Stage 1, 3-shift production schedule, using the min max end drying objective

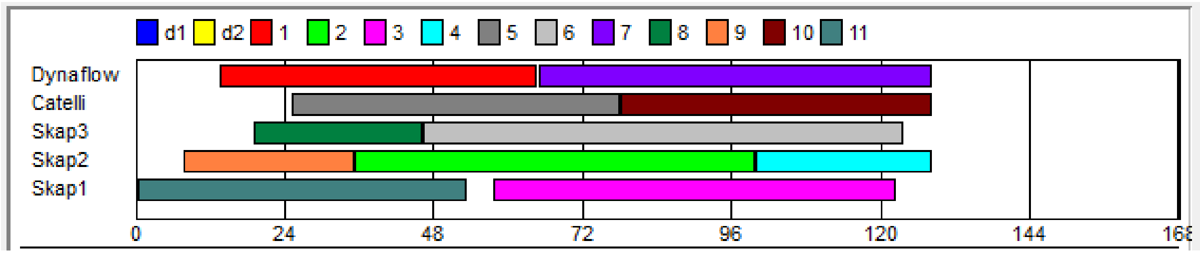

Figure 5.10: A Stage 2, 3-shift production schedule, using the min max end drying objective

We divide the figures in grids of 24 hours. The 168th hour reflects the end of the week when the following week is using a 3-shift schedule as well. We observe for example that Job 3 could start earlier in Stage 1, because Skap 1 is already free in Stage 2 and there is some time left in the schedule between Job 6 and Job 3 in Stage 1. This available space for adjustment has no effect on the objective value. We can easily do the adjustment for Job 3 manually after obtaining the schedule. Another option is scheduling more intermediates or more quantity, because if the schedule is using the weeks full capacity the available adjustment space does not occur. We give a graphical representation of the schedule from Week 5 using a 2-shift + Saturday schedule: in Figure 5.11 we present Stage 1 and in Figure 5.12 we present Stage 2.

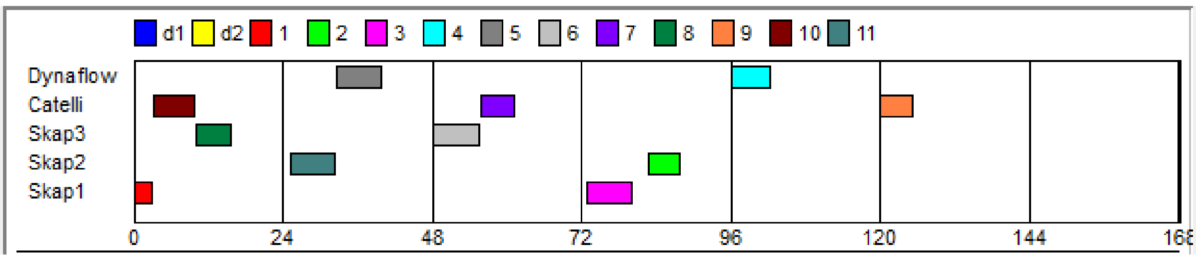

Figure 5.11: A Stage 1, 2-shift + Saturday production schedule, using the min max end drying objective

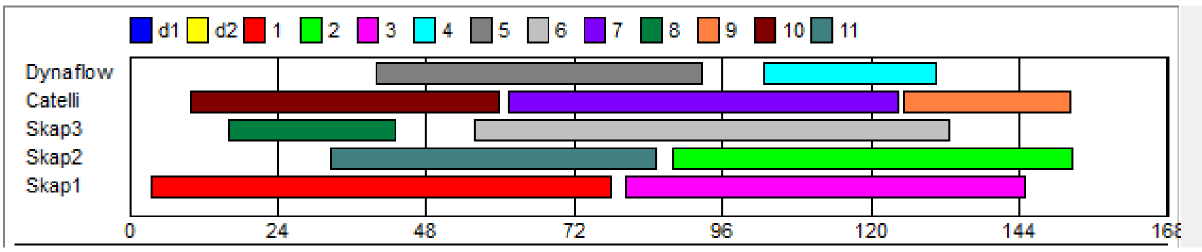

Figure 5.12: A Stage 2, 2-shift + Saturday production schedule, using the min max end drying objective

We again divide the figures in grids of 24 hours. The 168th hour reflects the end of the week and the start of the following 2-shift schedule. Using a 2-shift schedule, we can no longer schedule jobs in Stage 1 during the night. Therefore, production starts at each grid again, as this is when each production day starts, with 0 indicating Monday 6:30h, see Figure 5.11. When we schedule Week 5 using a 2-shift + Saturday schedule we observe that there is still space left for Job 2, 3, and 11 to increase their quantity, given that these quantity increases fit within the drying cabinet.

5.2 Case study Brynild: Heuristics

This section evaluates the experiments we conduct with the heuristics we base on our MILP. We describe our experiment design in Section 5.2.1. In Section 5.2.2 we summarize the results of the experiment design. Section 5.2.3 contains the evaluation by Brynild experts regarding the results and findings of our main model, and the heuristics.

5.2.1 Heuristic case study: Experiment Design

With our MILP we find the optimal schedule taking into consideration the pre-processed input data. We conduct experiments using heuristics, because finding the best solution requires multiple hours for most weeks. The goal of conducting the experiments with heuristics is to achieve a (near) optimal solution in less time, in minutes instead of hours, than with the main MILP.

We conduct experiments with the 3 heuristics we mention earlier in Chapter 4:

- Our MILP without sequence dependent set-up times;
- Our MILP that we decompose in an assignment and a sequencing MILP, which we compute sequentially;
- Our MILP with a maximum computation time of 10 minutes.

The instances for the experiments are the same 7 weeks as we use in Section 5.1. We test the 2 objectives for all 3 heuristics and obtain the objective values and computation time from the experiments. Next, we analyse the possibility of scaling down the shift schedules for every week using the heuristics. In total we conduct 24 experiments using 3-shift schedules, 36 experiments using 2-shift + Saturday schedules, and 6 experiments using 2-shift schedules.

When we compare the schedule we create using the heuristic without SDST to the main MILP, we incorporate the SDST. Because, in practice the SDST remain even when we do not consider them during scheduling. Therefore, we note that we schedule without taking into consideration the SDST, however we do add the required set-ups to the schedule before comparing the schedule to the main MILP.

5.2.2 Case study: Heuristic results

After describing the experiment design in Section 5.2.1 we evaluate the results. In Section 5.2.2.1 we start with analysing the results of the heuristics in comparison to our MILP, regarding the end-drying time objective. In Section 5.2.2.2 we do the same as in Section 5.2.2.1, however now for the Stage 1 objective. In Section 5.2.2.3 we compare the computation time of each heuristic to the computation time of our MILP. In Section 5.2.2.4 we analyse whether the heuristics are also able to find the schedules that we scale down in number of shifts.

5.2.2.1 Min max end-drying time for the heuristics

This section evaluates the end-drying time objective for the 3 heuristics in comparison to the MILP model, as we depict for each week in Figure 5.13.

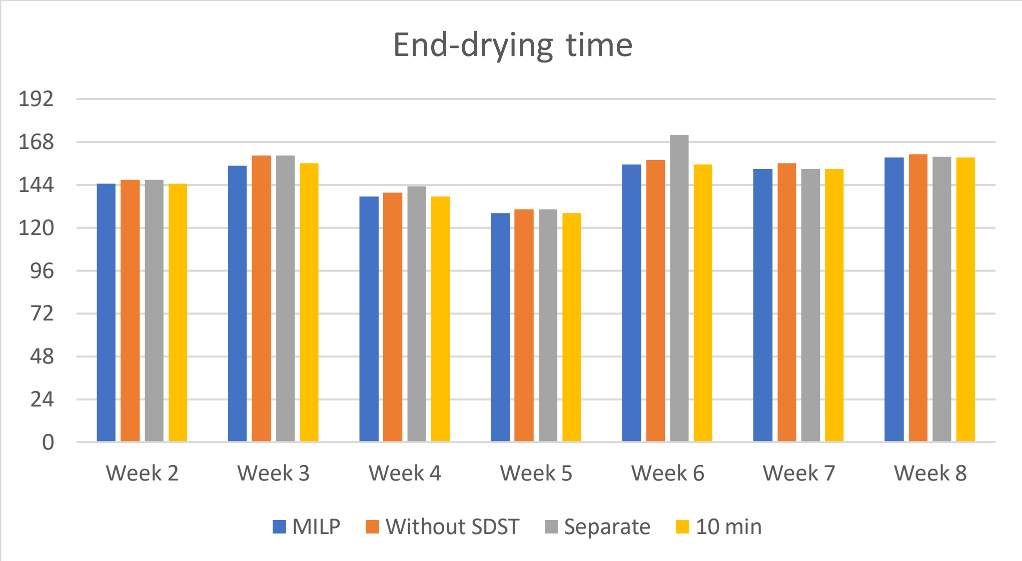

Figure 5.13: End-drying time for the heuristics and the MILP from Chapter 4

When we examine Figure 5.13 we observe that all heuristics perform relatively well. Especially stopping the MILP after 10 minutes gives good prospects. We analyse the difference of every heuristic in comparison to our MILP in percentages in Table 5.3.

Table 5.3: End-drying time heuristics compared to our MILP in percentages

End Drying	**Week 2**	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**	**Week 8**
Without SDST	+1.38%	+3.71%	+1.59%	+1.56%	+1.58%	+2.02%	+1.19%
Ass. & Seq. separate	+1.44%	+3.76%	+4.21%	+1.61%	+10.57%	0.00%	+0.33%
10 min run	0.00%	+0.87%	0.00%	0.00%	0.00%	0.00%	0.00%

Table 5.3 confirms that running the model for 10 minutes is an accurate heuristic. For almost all weeks the 10 min run heuristic finds the same objective value as our MILP. Only in Week 3, the heuristic finds a slightly less efficient schedule. Both the other heuristics find near optimal objective values, except for the 2-separate models in Week 6, where the objective is more than 10% higher. The difference between the weeks with a 2-shift + Saturday schedule, is that Week 6 includes more jobs with shorter drying times, and we assign all these jobs to the same drying cabinet. The objective for assigning the jobs, evenly distributes the drying time quantity over the drying cabinets and does not consider the number of jobs in 1 drying cabinet. Therefore, the reason regarding the 'bad' objective value could be that we assign these shorter drying times to the same drying cabinet in Week 6 as this is the main difference between weeks. The assignment of these shorter jobs to the same drying cabinet, in combination with the use of a 2-shift + Saturday schedule, can be the cause that the jobs are 'pushed' to the next day. If this 'push' happens on multiple days with multiple jobs in 1 drying cabinet, the end-drying time of the last job in that cabinet is relatively late. With separate models the jobs cannot interchange between drying cabinets when this occurs. We do not test this theory with other instances, but we expect that the 'bad' objective value in Week 6 for the heuristic that assigns and sequences the jobs separately is a combination of these factors. First, the assignment of short jobs to the same drying cabinet and second due to sequencing with production time windows, therefore the jobs are 'pushed' further in the schedule. Appendix F contains more information on the test results of the heuristics for the end-drying time objective.

5.2.2.2 Min max Stage 1 time for the heuristics

This section evaluates the Stage 1 end time objective for the 3 heuristics in comparison to the MILP model, as we depict for each week in Figure 5.14.

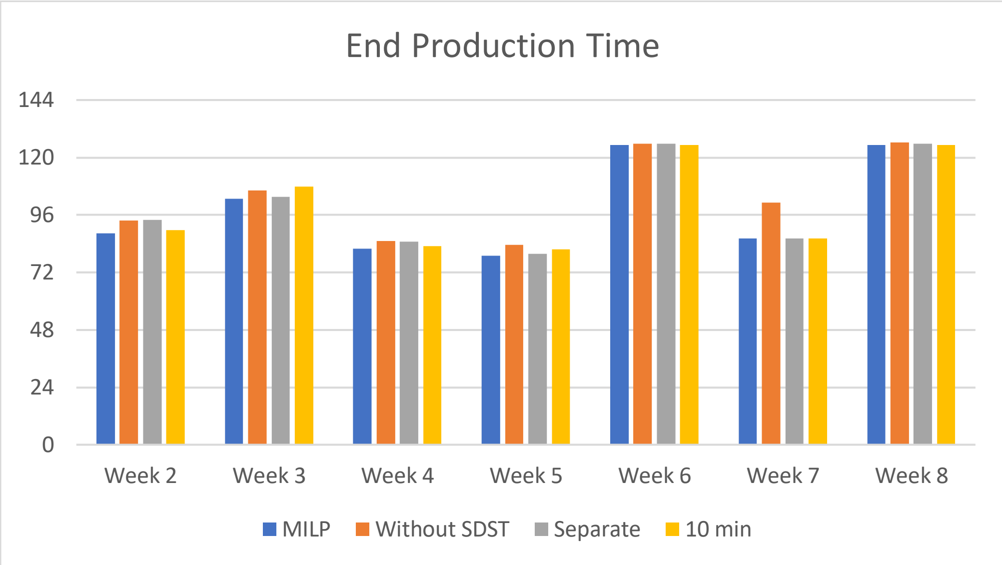

Figure 5.14: End time Stage 1 for the heuristics and the MILP from Chapter 4

We observe from Figure 5.14 that the heuristics have varying performances. For example, in Week 2, the schedule using the 2-separate models heuristic finds the worst objective value. In Week 3, the schedule using 10 min run heuristic, performs the worst, while in Week 7, the schedule using the model without SDST heuristic performs the worst. We also analyse the difference of every heuristic in comparison with the MILP in percentages, see Table 5.4.

Table 5.4: End time Stage 1 heuristics compared to our MILP in percentages

End Production	**Week 2**	**Week 3**	**Week 4**	**Week 5**	**Week 6**	**Week 7**	**Week 8**
Without SDST	+6.00%	+1.62%	+3.66%	+5.71%	+0.53%	+17.62%	+0.83%
Ass. & Seq. separate	+6.48%	+0.81%	+3.59%	+0.83%	+0.53%	0.00%	+0.42%
10 min run	+1.47%	+3.04%	+1.22%	+3.36%	0.00%	0.00%	0.00%

From Table 5.4 we observe that the 10 min run heuristic performs well for the 2-shift + Saturday schedules in weeks 6, 7, and 8. In weeks 2 to 5 however, the 10 min run heuristic performs the most stable with a variance of 0.826. We calculate the variance over the difference with the MILP. Therefore, it does not conclude how well a heuristic performs, but just how stable its performance deviation from the optimal solution is. The variance of the heuristic without SDST is 1.919 and for the 2 separate models the variance is 6.645.

The model without SDST has an outlier in Week 7; which is due to the 2-shift + Saturday schedule. When the set-up times are 2 hours instead of 0 or instead of 1.5 hours, the possibility arises that, with these longer set-up times we can only schedule 2 jobs, rather than the possibly previous 3 jobs per production day. We must schedule the additional job and the remaining jobs later because of this. We recognize a sort of bullwhip effect in this event. Even when we compare the practical model, without considering the SDST in the schedule, many set-up times remain 2 hours.

Appendix F presents additional information over the test results of the heuristics for the Stage 1 objective.

5.2.2.3 *Computation time of all heuristics and the main MILP for all weeks and both objectives*

We compare the objective results of the schedules we create using the heuristics, to the schedule we create using our MILP which we discuss in the previous section. In this section we compare the computation time of the heuristics to the MILP computation time for both objectives.

Due to time limitations for this research, we interrupt solving the models in some cases. We present the cases we interrupt in red. We stop the model when the integrality gap does not improve for more than 1 hour. Table 5.5 presents the computation time in seconds, for the Drying and Stage 1 objective.

Table 5.5: Computation time in seconds for the drying objective and Stage 1 objective

Drying **Week**	**MILP**	**Without SDST**	**Separate**	**10 min**	*Stage 1* **Week**	**MILP**	**Without SDST**	**Separate**	**10 min**
2	3840.97	21.56	2.14	600	2	12278.99	8075.22	3.34	600
3	4979.94	1316.81	26.31	600	3	42096.3	10314.06	9.73	600
4	147.48	8.83	2.81	147.48	4	12507.53	8196.91	2.34	600
5	1179.17	6.42	3.92	600	5	9940.61	6019.55	1.95	600
6	438.74	10.95	0.41	438.74	6	11426.4	31.63	0.78	600
7	379.91	6.16	0.37	379.91	7	466.94	7.61	0.62	466.94
8	498.31	13.42	2.82	498.31	8	6882.33	74.05	6.71	600

In Table 5.5 we display the computation times of the end-drying objective on the left part of the table and the computation times of the Stage 1 objective on the right part of the table. We only require interruption of models with the Stage 1 objective, for the MILP and for the MILP without SDST. The separate heuristic has the shortest computation time for both objectives. This is what we expect, because without the interaction between the assignment of parallel machines, and the sequencing of the single machine, the solution space becomes much smaller than the solution space for the other models. The main MILP and the MILP without SDST are solvable within reasonable time for the end-drying objective. However, for the Stage 1 objective, the weeks 2 to 5 take relatively long. Therefore, we conclude that models that consider both stages and have zero buffer capacity between these 2 stages, require more computation time to optimize Stage 1, than to optimize the whole schedule. This is opposed to models with infinite buffer capacity between the 2 stages, as in that case Stage 1 can optimize without considering the constraints from Stage 2. Then the model becomes a 'simple' sequencing problem, that requires less time than when the model must consider both stages.

All heuristics require the same amount or less time than our MILP. This is because the heuristics are simplifications of the MILP, and therefore require less time. We cannot state this with certainty for the heuristic MILP without SDST for the Stage 1 end time objective. Therefore, we evaluate the integrality gaps as well.

We analyse the integrality gap because we interrupt some models. When the integrality gap is small the model almost reached with certainty the optimal solution. When the integrality gap is still large, the probability is much higher that there exists a better solution than when the gap is small. We present the integrality gaps for the 10 min heuristic of the drying objective and the integrality gaps of the Stage 1 objectives in Table 5.6.

Table 5.6: Integrality gaps of the interrupted schedules

Drying		Stage 1			
Week	**10 min**	**Week**	**MILP**	**Without SDST**	**10 min**
2	1.53%	2	52.24%	24.42%	82.37%
3	27.04%	3	85%	27.17%	86.73%
4		4			33.39%
5	5.65%	5	68.98%	24.82%	82.98%
6		6			57.66%
7		7			
8		8	64.89%		74.39%

For both objectives, the integrality gaps are the largest for the 10 min heuristic in Week 3. Week 3 has the most jobs and is a 3-shift schedule, on top of that the 10 min heuristic stops the earliest. Therefore, the solution space is still the largest and the gap most likely requires the most time to completely close. For the Stage 1 objective, the integrality gaps are larger for the MILP than for the MILP without SDST. Therefore, we hypothesize that the MILP requires more additional time to complete the computation time. Moreover, all the MILPS we interfere, we stop later than the MILPs without SDST. Thus, we conclude that the SDST most likely finds their optimal solution faster than the MILP if we solve both to completion.

5.2.2.4 Reduction of shifts in schedule realized by heuristics

We test for every week whether we can create reduced schedules using the 3 heuristics. We present the results for this experiment in Table 5.7.

Table 5.7: Scaling down of the schedules using heuristics

Week	Model	Without SDST	Separate	10 min
2	Yes	No	Yes	Yes
3	No	No	No	No
4	Yes	No	Yes	Yes
5	Yes	Yes	Yes	Yes
6	No	No	No	No
7	Yes	Yes	Yes	Yes
8	No	No	No	No

The heuristic without SDST cannot find a reduced schedule for Week 2 and Week 4, whereas our MILP can find a reduced schedule. We assume that the reason that the heuristic cannot construct a schedule for these 2 weeks is the same type of bullwhip effect as we mention in Section 5.2.2.2. We therefore conclude that the sequence dependent set-up times are important to take into account when constructing the schedules. This applies particularly to schedules with high varying set-up times and multiple production time windows, such as the 2-shift schedules.

5.2.3 Case study: Evaluation by Brynild experts

The Brynild experts, supply chain manager and the planner/scheduler, are positive about the results. They agree that the objective to minimize latest end-drying time rather than minimizing the throughput for each intermediate individually, results in the desired schedules. The desired schedules are more flexible when dealing with unknown changes, reduce the constraints for upcoming weeks, and we can easily observe when the throughput can be higher.

The objective to minimize the end time of Stage 1 is of less importance to Brynild when scheduling 1 week ahead. Brynild states that they need to know the shift schedules 3 weeks in advance and that they assign staff members to complete shifts. Besides, Brynild wants to realise more throughput per week and for that goal, the end-drying time objective is of more relevance. The end-drying time objective requires less computation time and considers the next week as much as possible, since the end-drying time objective reduces the number of constraints for the upcoming week. Brynild experts agree that for an optimal schedule they do not mind running the model at night and obtaining the schedule in the morning. Furthermore, for testing various scenarios, Brynild wants to use the heuristic where the MILP stops after 10 minutes.

Following their positive reactions, the Brynild experts want to test the MILP with live data to see if they could use the MILP directly practice. We create a schedule for Week 7 and Week 8 of 2021 that we base on live data from Brynild. Although they are positive about the outcomes, there is one remark regarding the test schedule. In test Week 8, we start with the so-called Dent product. In practice however, Brynild does not prefer this product to start, since the materials for Dent products require pre-processing of 3 hours. Therefore, we add an additional constraint that ensures that the jobs containing Dent cannot start producing earlier than 3 hours after the start of the week. In practice it is also difficult to produce Dent products in the weekend, therefore we add a similar constraint for the weekends. The scheduler discusses the new schedule with the operators of the Støperi 1 team to collect feedback on whether the schedule can be applied in practice. The operators did not find any large discrepancy and positively comment that they '*find it hard to believe that the scheduling is done by a model'*.

In finetuning the schedules, the scheduler requests to change the input for the moulding/unloading to an average of 25 trays per minute, instead of 27 trays. Furthermore, they like to incorporate the full drying cabinets from the week before, for when the drying does not finish before the beginning of the week. We can easily add constraints for the specific drying cabinet that the model cannot schedule them before a certain time. After incorporating these remarks, we schedule Week 9 with live data. Unfortunately, Brynild momentarily has a lot going on in the factory and they could not find the time to analyse Week 9.

5.3 GENERIC USE OF THE MATHEMATICAL MODEL

In this section we evaluate the experiments regarding the general capabilities of our MILP. In Section 5.3.1 we present experiment designs, where we explain the reasoning behind conducting certain experiments and what we examine. In Section 5.3.2 we present the results of the experiments. During these experiments we solely use the end-drying time objective.

5.3.1 General capabilities: Experiment design

In Section 5.3.1.1 we present the experiment design for examining the computation time limits of our MILP, testing with various numbers of jobs and parallel machines. In Section 5.3.1.2 we present the experiment design to examine multiple variants of our MILP without buffer capacity. We base both experiments in this section on different job instances:

- Small jobs: short drying time on the parallel machines;
- Large jobs: long drying time on the parallel machines;
- Mixture of small and large jobs.

In Appendix E, we give the details regarding how we construct these job instances that we base on Brynild's input data.

5.3.1.1 *Experiment design: Computation time limits of our main MILP*

The goal of this experiment is to indicate for what size, in terms of number of jobs and number of parallel machines, we can compute the entire model in 1 night, that is 16 hours. With the results of the experiments, we can approximate whether we can compute future model sizes within a reasonable amount of time, or if we should recommend using a heuristic to solve the problem instead. Next, we analyse if the various types of jobs, e.g., small, large, or a mixture of both have an influence on the computation time.

During this experiment we test the 3 job types for 2, 6 and 10 parallel machines. We start with 3 jobs and we increase the number of jobs until the computation time requires more than 16 hours (960 minutes). When the job instances are too large and require more than 16 hours, we assess the integrality gap. We mainly focus on the increase of jobs as in the Brynild case study, we learn that the difference between the number of jobs has a large effect on the computation time of the model. Our other focus is on the number of parallel machines as these largely influence the size of the solution space.

To conclude, we conduct experiments using 3 different job types, and for 3 different numbers of parallel machines. The number of jobs we test depends on when we reach a computation time that requires more than 16 hours. The experiments give a general idea of when the instances are too large to compute within 16 hours. We must note that a better processor could theoretically solve the instances faster.

5.3.1.2 *Experiment design: 8 models with zero buffer capacity of our MILP*

The goal of this experiment is to enhance the literature on the topic of models with zero buffer capacity. The first experiment is to run the model for 5 minutes. The goal of this experiment is to observe whether a model reaches optimality within 5 minutes, or if an integrality gap remains. In the second experiment we let the model run until the integrality gap is zero and the model creates the optimal solution. Thereafter, we compare the objective value after 5 minutes to the optimal objective.

In Chapter 4 we describe a two-stage Hybrid Flow Shop with non-identical parallel machines, sequence dependent set-up times, production time windows, and zero buffer capacity. The newfound idea of this MILP is the absence of buffer capacity between Stage 1 and Stage 2 of the HFS. To broaden our scope on the zero buffer capacity subject, we conduct the experiments using various mathematical models without buffer points. We conduct the experiments using 8 different MILPs, see Table 5.8.

Table 5.8: Mathematical models without buffer capacity

Model	**Non-identical parallel machines**	**Sequence dependent set-up times**	**Production time windows**
1	X	X	X
2	X	X	
3	X		
4		X	
5			X
6		X	X
7	X		X
8			

We present the mathematical description of the models from Table 5.8 in Appendix D. Model 1 is the main model we describe in Section 4.2. Model 2 has many commonalities with the 3-shift schedules of the Brynild case study, with the only difference that in this model unlimited production hours are available. Model 7 is a familiar model as well; this model is one of the heuristics we review during the Case study in the Section 5.2.

We conduct 2 experiments with the 8 models out of Table 5.8. We test 3 different instances for these 8 models. The instances consist of 5 parallel machines and 10 small jobs, large jobs, or a mix of both. We specifically use this composition of instances, as we find for the experiments from Section 5.3.1.1 that, for the main MILP, these instances probably have a computation time under 16 hours. The 3 instances are the same for all 8 models, to allow a comparison between the models. When there is no SDST, the set-up time is 2 hours. The production time windows are 16 hours each. The identical parallel machines are able to produce all jobs, and all cabinets have the processing time of the fastest drying cabinet.

5.3.2 General capabilities: Results

In Section 5.3.2.1 we present the results for the experiments concerning computation time limits of our model. In Section 5.3.2.2 we present the results of the 8 models that have zero buffer capacity between the 2 stages.

5.3.2.1 Results: Computation time limits of our main MILP

We present the computation time in minutes for the instances including jobs with short processing times in Figure 5.15.

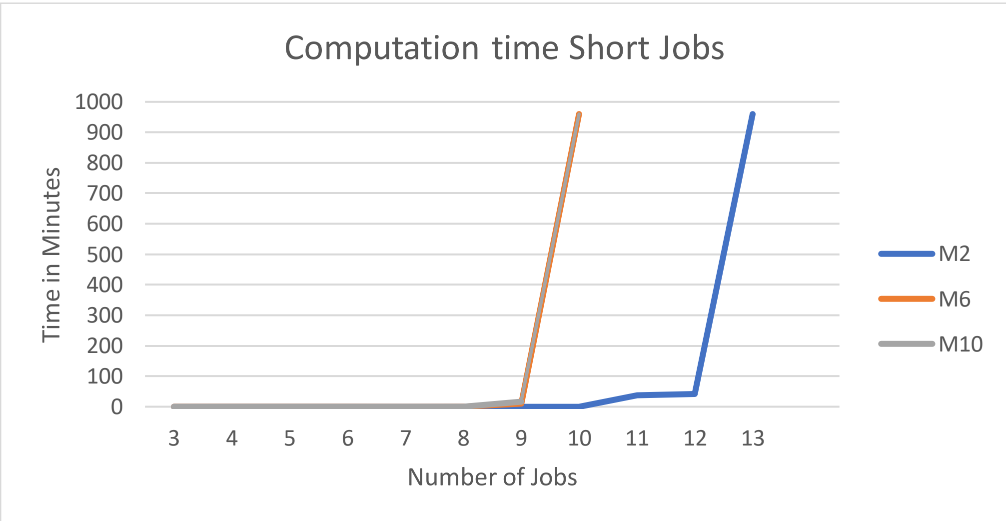

Figure 5.15: Computation time for short jobs in minutes for 2, 6, and 10 parallel machines

Figure 5.15 indicates that the computation time for the jobs with short processing times abruptly increases as soon as we add one job (too many). We can include, and compute within a reasonable amount of time, more jobs for the instance with 2 parallel machines than the instances with 6 and 10 parallel machines. Both instances with 6 and 10 parallel machines can easily create a schedule for 9 jobs, however not for 10 jobs. The integrality gap for 6 parallel machines is 16.3%, and the integrality gap for 10 parallel machines is 19.82%. Considering these integrality gaps, we assume that the instance with 6 parallel machines is closer to optimality than the instance with the 10 parallel machines. Therefore, we conclude for short jobs, that more parallel machines lead to less possible jobs to schedule within a reasonable amount of time.

We present the computation time in minutes for instances including jobs with long processing times in Figure 5.16.

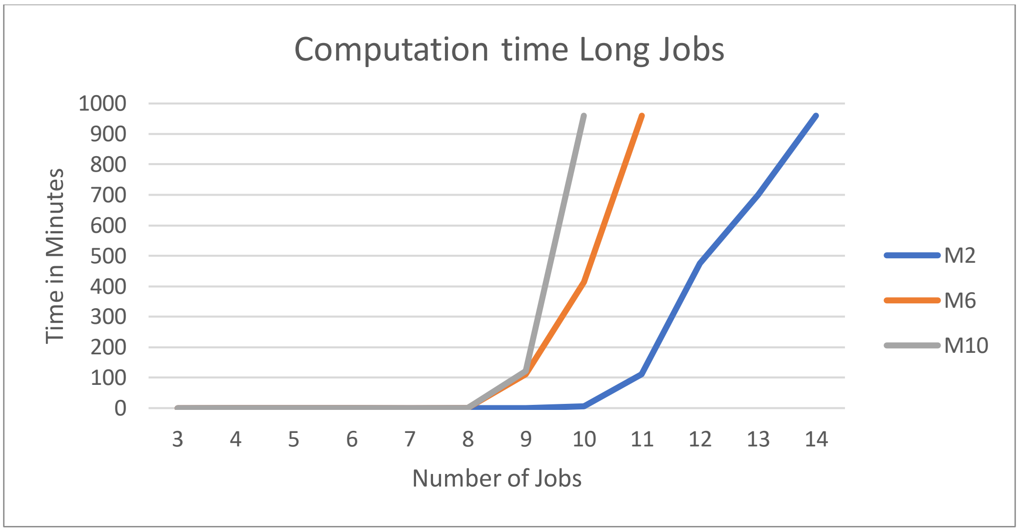

Figure 5.16: Computation time for long jobs in minutes for 2, 6, and 10 parallel machines

Figure 5.16 confirms our findings regarding Figure 5.15, that the more parallel machines we include, the less jobs we can schedule in a reasonable amount of time. The computation time for 2, 6, 10 parallel machines gradually increase for the instances where jobs contain long processing times. We observe that the computation time is for some instances higher than for the short jobs, as for example we see in the case of 2 parallel machines. We suspect that the computation time is longer for certain instances because we must use more production time windows to compute the jobs with longer processing times. These time windows influence the computation time as well.

We present the computation time in minutes for instances including jobs with mixed processing times in Figure 5.17.

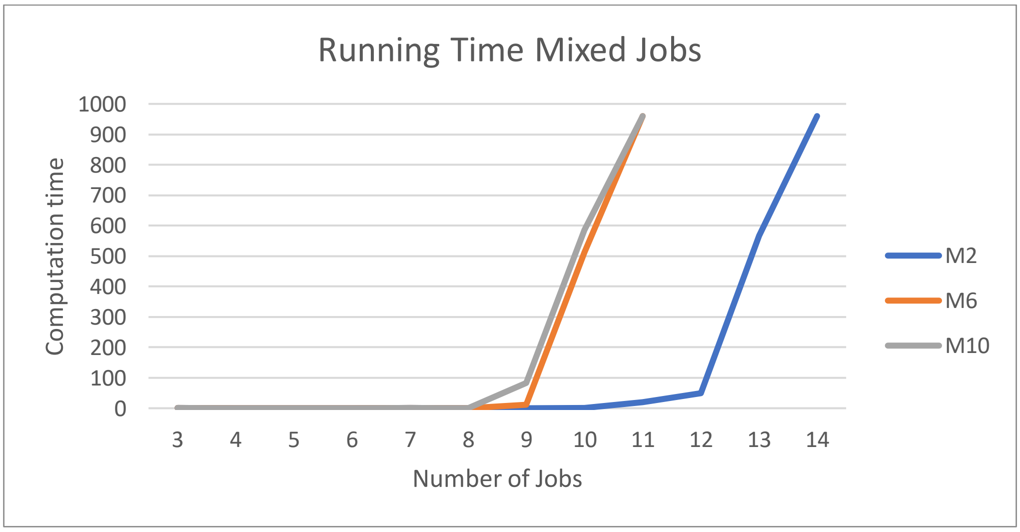

Figure 5.17: Computation time for mixed jobs in minutes for 2, 6, and 10 parallel machines

Figure 5.17 is a combination of Figure 5.15, where the computation times increase abruptly, and of Figure 5.16, where the computation times increase gradually. We can explain this combination by the fact that the instances contain both jobs with long and short processing times.

For all instances with different types of jobs and the use of various number of parallel machines, we observe an exponential increase in computation time when the number of jobs increases. Besides, the more parallel machines per instance, the less jobs we can schedule within a reasonable amount of time. Another observation is that the difference in number of jobs is higher when we vary between 2 and 6 parallel machines, than the difference in jobs between the use of 6 and 10 parallel machines.

5.3.2.2 Results: 8 models with zero buffer capacity of our MILP

After conducting the experiments for the 8 models, we evaluate the objective values of each model, and for each instance, with the objective value we find after 5 minutes. The largest difference we observe is for Model 4 in Instance 1, with a difference of 3.05%. For more than 70% of the outcomes we find the same objective value within 5 minutes computation time, as after completion of the full computation time. Only Model 1 has a small difference for every instance; 2.98% for Instance 1, 0.79% for Instance 2, and 0.22% for Instance 3. Therefore, we conclude that for these instances' sizes, most models find the optimal solution within 5 minutes. We present the comparison of the objective value of the optimal schedule to the objective value after 5 minutes graphically in Appendix F.

We compare the objective (diagram) and the final computation time (line) of the 3 instances for all 8 models in Figure 5.18.

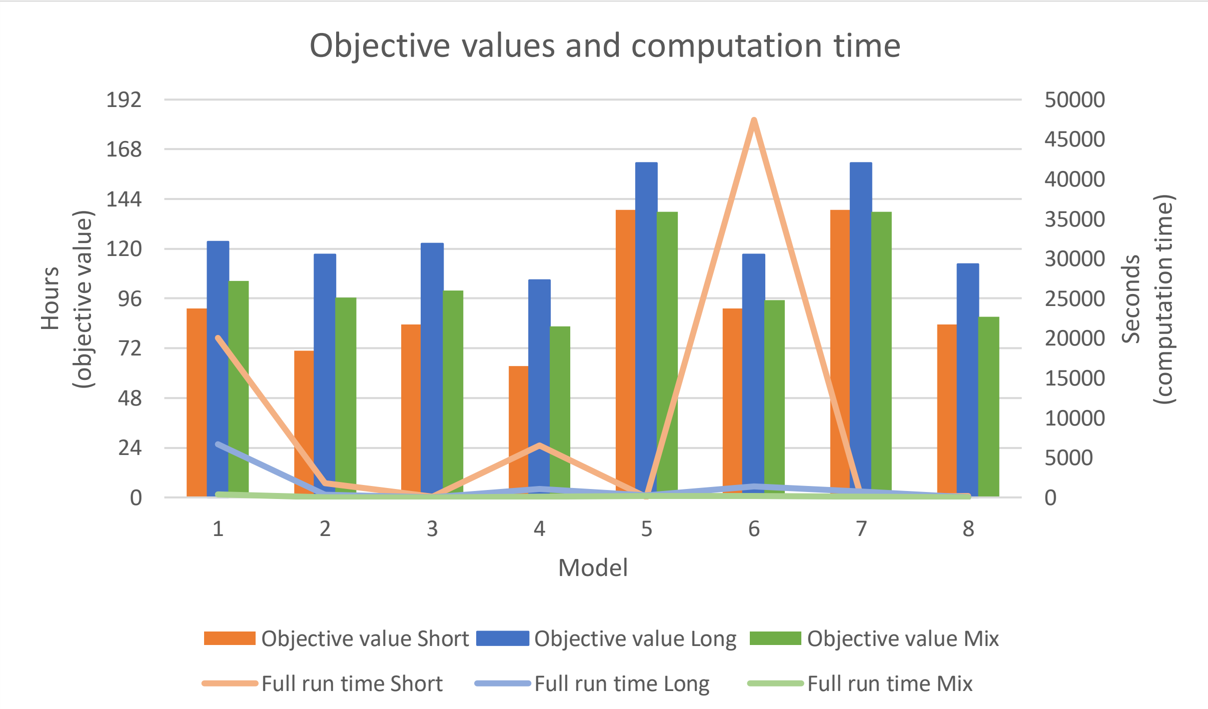

Figure 5.18: Comparison of the 8 models for 3 instances in terms of objective value (diagram) and computation time (line)

Note that the computation time for Instance 1, with short drying times, is for almost all models the highest; this is in line with the patterns we observe in Section 5.3.2.1.

The objective values are in line with what we expect:

- Higher objective values for models with production time windows;
- Lower objective values for models with SDST;
- Higher objective values for models with non-identical parallel machines.

We explain the lower objective values for models with identical parallel machines by the fact that the identical parallel machines have the shortest drying time out of the two available drying times. Moreover, all machines can produce all jobs. Thus, shorter processing times and a larger solution space lead to a higher probability of obtaining a lower objective value. The lower objective for models with SDST is explainable since the set-up times become 2 hours without SDST. These changes in objective values arise as a result of how we conduct the instances. For example, if we state that the models without SDST have zero set-up time, the objectives of models without SDST are lower instead of higher.

Model 5 and Model 7 have the highest objective of all instances. Model 5 and Model 7 have in common that there are no sequence dependent set-up times in place, and that there are production time window constraints. The other models do not share the combination of these 2 features. Therefore, it is reasonable to conclude that in our case the combination of no SDST and the presence of production time windows, leads to the highest objective values.

The objectives depend heavily on what values we assign to the parameters of the different models. Therefore, we zoom in on the computation time of the various models. The values of the parameters have less influence, however the amount of various parameter values does have considerable influence on the computation time. As the number of values for parameters is the same for all models, and since what parameters we include depends on the models, we make a comparison that is mainly influenced by the model type.

We isolate every character trade by comparing 2 models that differ in only one of the 3-character trades. For example, comparison of Model 2 to Model 3 to evaluate the effect of SDST on the computation time. We present our evaluation of the findings in Table 5.9, Table 5.10, and Table 5.11.

Table 5.9: Difference in percentage in computation time when there are no SDST

SDST	**Model 1 vs Model 7**	**Model 2 vs Model 3**	**Model 4 vs Model 8**	**Model 5 vs Model 6**
Instance 1	-99%	-94%	-98%	-100%
Instance 2	-89%	-85%	-98%	-81%
Instance 3	-79%	-97%	-98%	-6%

Table 5.9 indicates that for every model and for every instance, the model without the sequence dependent set-up times computes much faster than the model with SDST. Only the comparison between Model 5 and Model 6 for Instance 3 gives a small difference in computation time. We conclude that all models without the SDST compute faster, like we in Section 5.2 partly observe with the heuristic as well.

Table 5.10: Difference in percentage in computation time when there are no time windows

Time Windows	**Model 1 vs Model 2**	**Model 3 vs Model 7**	**Model 4 vs Model 6**	**Model 5 vs Model 8**
Instance 1	-91%	-17%	-86%	262%
Instance 2	-95%	-93%	-20%	-92%
Instance 3	-85%	-98%	-30%	-99%

Table 5.10 indicates for almost all models, except for Model 5 in comparison with Model 8 for Instance 1, that not taking time windows into consideration, decreases computation time. Model 5 does not consider non-identical parallel machines, and SDST, but does consider production time windows. Model 8 is the basic model, which does not consider any of the 3 characteristics. We do not know why Model 5 computes faster than Model 8 for the instance with jobs containing short processing times. It could be a coincidence that we compute Model 5 for this specific instance very fast. Therefore, we test the assumption for other instances with jobs containing short processing times. Model 5 requires

much longer to solve than Model 8 for these other instances with jobs containing short processing times. We do not know the specific reason why this specific instance 1 solves faster for Model 5.

Table 5.11: Difference in percentage in computation time when there are no identical parallel machines

Non identical parallel machines	**Model 2 vs Model 4**	**Model 5 vs Model 7**	**Model 1 vs Model 6**	**Model 3 vs Model 8**
Instance 1	271%	-69%	137%	33%
Instance 2	231%	-66%	-80%	-61%
Instance 3	134%	118%	-51%	19%

In Table 5.11 we observe that computation time increases and decreases for models with identical parallel machines. We explain this increasing and decreasing in computation time by the fact that the non-identical parallel machines differ in two ways. Firstly, there are different processing times for different parallel machines. Secondly, not all parallel machines can dry all jobs. If changing the first, that is all parallel machines have the same processing time, then in theory the model should compute faster. However, when changing the second, that is all parallel machines can dry every job, this makes the solution space larger and therefore, in theory the model requires more computation time. The 2 differences between non-identical parallel machines and identical parallel machines have different weights in different situations, as we see in Table 5.11. We confirm this hypothesis by changing only one of the two differences, and we obtain the effect we expect. Thus, shorter computation times with every parallel machine having the same drying time, and longer computation times in the case that all the parallel machines can dry every job. We conclude that the models without non-identical parallel machines solves faster or slower depending on the situation. It is important to know the difference between the identical and non-identical parallel machines. We do not analyse why a specific characteristic is more influential than the other in the different situations.

5.4 CONCLUSION

This chapter discusses the Brynild case, where we evaluate the experiment design and results of the main MILP and the heuristics. Next, we evaluate the general characteristics of our MILP and other models, which similar to our MILP, do not have any buffer capacity between 2 stages of the HFS.

5.4.1 Brynild case

We examine 7 weeks, namely Week 2 to Week 8 of 2020. Originally Brynild schedules the first four of the seven weeks using a 3-shift schedule, and the remaining 3 weeks using a 2-shift + Saturday schedule. We evaluate different heuristics in the pre-processing phase. The heuristics regard the dividing of products over multiple jobs, in the case that the quantity of the product is too large to fit in 1 drying cabinet. After evaluating the results, we decide to use the same rule as Brynild originally does: divide the quantity evenly over the minimum number of jobs.

We evaluate 2 objectives. The first objective is the drying objective; the minimization of the latest end-drying time. In all the schedules we create the drying finishes significantly earlier ($p<0.05$). On average 38 hours earlier than in Brynild's original schedules. The second objective is the Stage 1 objective; the minimization of the maximum end time in Stage 1. In all schedules, Stage 1 finishes significantly earlier ($p<0.05$), on average 17.5 hours earlier than in Brynild's original schedules. Next, we observe that the set-up time does not have a direct effect on the end time of Stage 1. Therefore, we conclude that the availability of the drying cabinets is likely more influential to the Stage 1 end time than the set-up times are.

We can scale down four out of the seven schedules. We can schedule three out of four original 3-shift schedules using a 2-shift + Saturday schedule, and one of the three original 2-shift + Saturday schedules using a 2-shift schedule.

Out of the 3 heuristics we mention in Chapter 4, the heuristic where we interrupt our MILP after 10 minutes, performs the best considering the objective values and the computation times of both objectives. Moreover, we find that considering sequence dependent set-up time is very important for creating schedules. Especially when production time windows are in place.

5.4.2 General

To examine the limits of our MILP we research 3 types of jobs: with short, with long, and with a mix of short and long drying times. Furthermore, we test different numbers of parallel machines. For all 3 types of instances and for the different number of parallel machines, we observe an exponential increase in computation time when the number of jobs increase. Next, the more parallel machines per instance, the less jobs we can schedule within 16 hours. For the short jobs we observe a sudden increase in computation time, while for the long jobs and the mixed jobs the increase in computation time develops more gradually.

We examine 8 two-stage models varying in: including production time windows, SDST, and non-identical parallel machines, that share the characteristic of zero buffer capacity between the 2 stages. We conclude that models without sequence dependent set-up times compute much faster than models with sequence dependent set-up times. Besides, in almost all cases the models without production time windows, compute faster than the models with production time windows. Lastly, it variates whether models with identical parallel machines compute faster than models without identical parallel machines. After further evaluation, we hypothesize that when solely the process time becomes the same for the parallel machines, the model computes faster, and when all parallel machines are able to process certain jobs, the solutions space becomes larger and the computation times higher.

6 CONCLUSION & RECOMMENDATIONS

Chapter 6 concludes our research; we give short answers to every research question in Section 6.1. In Section 6.2 we discuss Brynild's case study. In Section 6.3 we discuss the scientific value of our work. Lastly, in Section 6.4 we give recommendations on how to move forward with implementing the MILP.

6.1 RESEARCH QUESTION CONCLUSIONS

This section answers the research questions as we formulate in Chapter 1. Together, these sub research questions answer the main research goal.

Research question 1: *What does the current situation at Brynild look like?*
Brynild's production line consists of 3 parts; Støperi 1, Drage Sukker, and Godteri Pakking. The biggest bottleneck is the drying area, which is part of Støperi 1. The production scheduling of Støperi 1 is done manually and experience-based. We formulate the scheduling problem of Støperi 1 as: *a problem with non-identical parallel machines, sequence depending set-up times, production time windows, and without buffers.* The data analyses of the current scheduling methods indicates that the bottleneck, the drying area, is almost always occupied (89% to 98% of the time), but not fully utilized (40% to 57% of the time).

Research question 2: *Which methods are described in literature regarding the scheduling of a production line similar to Brynild's?*
We identify Brynilds Støperi 1 line as a two-stage hybrid flow shop with non-identical parallel machines, sequence dependent set-up times, production time windows, and without intermediate buffer points. This exact problem is not known in literature. Therefore, we evaluate 6 problems with similar characteristics. It appears that these similar models are solved using a meta-heuristic such as tabu search, or genetic algorithms. In Brynild's case however, there is no buffer capacity between the 2 stages, and production time windows are in place. Due to the combination of these 2 conditions is it difficult to obtain good solutions using integrated meta-heuristics, therefore we prefer not to use them to compute our HFS problem. Luckily, the instances of Brynild's problem are likely small enough to compute within reasonable time using a Mixed Integer Linear Programming model.

Research question 3: *How can we develop a scheduling model that improves the throughput of the production line?*
Before developing a scheduling model that improves the throughput of the production line, we assume that we structure the input data in such a way that jobs consist of 1 or more intermediates, which can dry together and, in terms of quantity always fit in 1 drying cabinet. Another assumption we do is that the drying cabinets have the same capacity, and are available at the beginning of the week. We develop the main MILP model that considers a two-stage hybrid flow shop in which the drying cabinets are non-identical parallel machines. The model includes sequence dependent set-up times, production time windows, and there is zero buffer capacity between the 2 stages. The objective of the model is to minimize the make span of the schedule. Next, we introduce 3 heuristics that we base on our main MILP:

- Our MILP without sequence dependent set-up times;
- Our MILP that we decompose in an assignment and a sequencing MILP, which we compute sequentially;
- Our MILP with a maximum computation time of 10 minutes.

Research question 4: *How does the proposed scheduling method perform?*
We examine 7 weeks: Week 2 to Week 8 of 2020. The first four of the seven weeks consist of a 3-shift schedule and the remaining three weeks of a 2-shift + Saturday schedule. We evaluate 2 objectives. The first objective minimizes the make span up to the end-drying time, and the second objective the make span up to Stage 1. Both objectives provide significantly better schedules in comparison to Brynild's original schedules. Using the first objective we create solutions that on average finish 38 hours earlier, while with the second objective finishes on average 17.5 hours earlier. We could schedule four of the seven weeks using a shift schedule that requires less shifts. The heuristic that interrupts the MILP after 10 minutes performs the best. For the end-drying time objective, the heuristic obtains the same objective value for six of the seven weeks, and in the other week an objective that is 0.87% higher. Brynild agrees that the scheduling method is very representable for practice and that the scheduling methods only needs a few small changes before Brynild could use the method in practice.

We evaluate 2 generic experiment designs. The first experiment is to evaluate the limits of our MILP in terms of computation time. We observe an exponential increase in computation time for all 3 types of instances (short, long, or mixed processing times), when the number of jobs increase. The second experiment is to evaluate 8 two-stage models varying in rather or not including production time windows, SDST, and non-identical parallel machines, that share the characteristic of having zero buffer capacity between the 2 stages. We conclude that models without sequence dependent set-up times compute much faster than models with sequence dependent set-up times. Next, in almost all cases the models without production time windows compute faster than the models with production time windows. Finally, a model including identical parallel machines shows variating results in whether the computation requires less or more time. After further evaluation, we conclude that when solely the process time becomes the same for all parallel machines, the model computes faster. When more parallel machines are able to process certain jobs, the solutions space becomes larger and the computation time higher.

Concluding, our proposed MILP model, supplemented with the 10 minute run heuristic, is a good scheduling method that can successfully be used by Brynild. Brynild confirms that the scheduling method would be able to realize higher throughput of the production time and that it is feasible to use in practice. Therefore, our MILP model realizes the research goal, which is: *To develop a scheduling method that improves the scheduling of the production orders under consideration of multiple process constraints*.

6.2 DISCUSSION BRYNILD CASE STUDY

In this discussion we evaluate the largest limitation of our model, which is of interest for Brynild. Next, we discuss multiple improvement suggestions for the model. Thereafter, we discuss further research areas for Brynild to explore in order to improve the scheduling of their production line. Lastly, we discuss Brynild's future in terms of expanding the production line and using our scheduling method.

The largest limitation of the model of importance to Brynild is that in our MILP we consider 5 drying cabinets all with the same capacity. While in practice there are 7 drying cabinets; Dynaflow and Catelli consist of two times 2 drying lanes, instead of one time 4 drying lanes. By considering solely 5 drying cabinets we lose scheduling flexibility. However, if we consider 7 drying cabinets, we need to acknowledge 1 additional characteristic for the drying cabinets, namely the capacity of each drying cabinet. Due to this additional characteristic, the problem becomes larger. The model needs to be able to dry jobs together, since all jobs obtain the maximum size of the smallest drying cabinet. If we construct the MILP in this way, the instances contain 7 drying cabinets, and almost two times the number of jobs. When observing the results of the experiments regarding the computation time limits, we conclude that these instances do not compute within the reasonable amount of time of 16 hours.

Another way to incorporate the different sizes of drying cabinets is by using a MILP that considers batching. The disadvantage of batching is that the theoretical model is very difficult to produce in practice. The trade-off Brynild can make, is to assume 5 drying cabinets and obtain an optimal schedule for these 5 drying cabinets. Or alternatively, to incorporate jobs that dry together and schedule 7 drying cabinets in order to obtain a (near) optimal solution with a heuristic.

We have several improvement suggestions for Brynild to make the practical use of the MILP simpler. The first suggestion makes the practice use of the MILP simpler if Brynild wants to use the scheduling model for unanticipated changes in the production planning as well. We suggest using a rolling horizon instead of a finite horizon. With a rolling horizon the schedule can easily adapt when unexpected changes arise. We also suggest to assign a priority to the remaining jobs if changes occur. We need the priority when it appears that we cannot produce all jobs in one week any longer. The priority indicates, which jobs we need to produce in that specific week and, which jobs we 'only' prefer to produce in that week. We can include the priority in the objective or as a constraint.

Brynild indicates that they prefer to schedule maintenance as well. When the maintenance still requires planning we suggest that Brynild defines a specific job representing the maintenance time. After running the MILP, the model schedules the maintenance on the best time possible for Brynild. In the case that the maintenance is already planned at Stage 1, we suggest adjusting the production time windows of the model. When maintenance is planned for a drying cabinet, we need to add an additional constraint, which indicates that we cannot use that specific drying cabinet during the maintenance period.

We have another suggestion regarding the practical use of the scheduling method. Brynild should use the objective to minimize the end time of Stage 1 to create a week schedule 3 weeks in advance, and to determine the minimal number of shifts. Brynild can use this information to select the shift schedule. Brynild must take into account that certain changes can occur and should therefore incorporate some slack for the number of shifts. Accordingly, we suggest that Brynild uses the MILP with the end-drying objective, in which they incorporate the selected shift schedule as a constraint.

We have various suggestions for further research that can be beneficial for the scheduling of the production line from Brynild's confectionary. The first recommendation is to examine the schedule for multiple weeks, using the end time Stage 1 objective and add as constraint that the drying cabinets need to finish drying before the start of the upcoming week. In our 7 weeks from 2020 we observe much potential for this suggestion. All drying cabinets finish drying in time and we can minimize the Stage 1 end time. However, it can occur that the model cannot find a feasible schedule. For example, for the 3 weeks we schedule with the end-drying objective using live data from Brynild, two of the three schedules feature drying cabinets that finish drying in the first shift of the following week. With a constraint on the end-drying time, these instances would be infeasible to schedule. Therefore, we suggest exploring the probability that these types of seemingly infeasible instances occur, before using the model in this way.

The second suggestion is to research various heuristics dividing the quantity over jobs in the pre-processing phase. Heuristic rule 2, fill the jobs completely and the last job contains the rest quantity, has much potential to create better objectives values. Brynild could test different input instances and may find patterns. In this study, alternative job configurations should be taken into consideration as well. Thus, not solely those heuristics that lead to the minimum number of jobs. This research can provide suggestions for different situations regarding when to use what heuristic. Brynild can use the 10 minuntes' heuristic to research whether different job configurations can be interesting for obtaining better schedules. We expect that using the correct heuristic in a particular situation can reduce the objective with 8 or more hours, like we observe in Week 8 during our limited research regarding heuristics.

The third research suggestion has an influence on the parameters of our MILP. Brynild indicates a variability in production volumes. Therefore while scheduling, 2 pallets need to remain empty. Brynild can obtain useful information when they further examine this variability for the various products. Thereafter, Brynild can do a trade-off regarding the 2-pallet rule for each product. In this trade-off, Brynild can weight the costs of materials against the profit resulting from the additional drying space. Before making the trade-off, we should know what the variability of each product is, as this could be different for every product. Brynild can conduct a stochastic simulation, using Monte Carlo, to obtain the insight they need for this trade-off. The Monte Carlo can also help to determine for every intermediate the number of trays or pallets that Brynild should schedule empty or maybe the number of trays or pallets that Brynild should schedule to overproduce. When this information is available, the theoretical schedule represents practice better, and the probability for obtaining more profit increases.

A last recommendation for Brynild to further research is the possibility of scheduling more than one week at a time. We assume a 1-week scheduling horizon since the output of the planning is cost-based. Interchanging intermediates would nullify the planning output. Therefore, if for example Brynild would become interested in a scheduling horizon of 2 weeks, the planning should give weights of importance to the production of specific intermediates in certain weeks. Brynild can incorporate these weights in the objective or as a constraint. Changing the scheduling horizon to 2 or even to 3 weeks can improve Brynild's throughput as the solution space becomes much larger.

The future of Brynild might involve expanding the complexity of their production line by adding an additional drying cabinet. Furthermore, the number of jobs is likely to increase as well. As Brynild is currently at the limit of scheduling within 16 hours (with approximately 11-12 jobs per week) with the use of our MILP model, there is a high probability that they cannot schedule all week instances to optimality within 16 hours. In this case, we recommend Brynild to still use the MILP, but interrupt the MILP after these 16 hours, as the (near) optimal solution is most likely already found.

6.3 Discussion to Further Enhance the Literature Value

In this discussion we evaluate the limitations of our model, which are of interest to the literature. Next, we recommend further research to enhance the literature.

The limitation with the biggest impact on our MILP, is the assumption of pre-processing the input data. We structure the input data in such a way that jobs consist of 1 or more intermediates, which can dry together and, in terms of quantity always fit in 1 drying cabinet. This is a reasonable assumption for the Brynild case as in this way the schedule becomes more useful in practice. However, the general experiments we conduct with our MILP represent only part of a two-stage hybrid flow shop problem without buffers. Therefore, we recommend researching other MILPs including batching and different capacities for the non-identical parallel machines. As the parallel machines differ regarding 3 characteristics, we recommend testing every characteristic individually and in different combinations in order to evaluate the effect of each of the characteristics. We also recommend conducting the same experiment design for the MILP that considers batching, like we do for the 8 different MILPs. Our experiments can be used as a benchmark to evaluate the influence of batching on the MILPs' computation time and objective values. We hypothesize that the model with batching finds better objective values, however also that the computation requires much more time.

During the evaluation of our experiments, some questions remain unanswered. One of these is "*why is Model 5 with production time windows solved much faster than Model 8 without production time windows, for Instance 1*". We find that this is not the case for other instances with short jobs. Therefore, there is a high probability that this discrepancy is specific for this instance and an exception to the rule. In this stage of our research, we recommend to not focus on these very specific cases.

When more of these discrepancies arise, we could conduct a research to explore underlying causes. For now however, we recommend focussing on a broader aspect, including the zero buffer capacity between stages.

During the research regarding the limitations of our MILP, we solely differentiate between instance types, number of parallel machines and the number of jobs. However, also other parameters have impact on the computation time. For example, the number of different set-up times. We recommend to explore all these different parameters and to evaluate, which parameters have the most influence on the computation time of the model. We assume that the parameters we already tested are the ones with the most influence, since that is why we tested them. Besides, we observe a large difference between the number of jobs that we can schedule for 2 parallel machines and for 6 parallel machines within 16 hours. Therefore, it could be interesting to evaluate 3, 4, and 5 parallel machines as well. Another interesting observation is that the computation time of the jobs with short processing times increases abruptly, and that the computation time for jobs with long processing times increases gradually. We suspect that this has to do with the number of production time windows. However, further research is required to reinforce this suspicion or to reject the hypothesis and find another explanation.

We further recommend testing more (different types of) instances for the 8 different MILPs. With more data, the differences can be quantified as averages. Consequently, we may then be able to state for example: when characteristic Y is not considered, on average the model computes X% faster. The specific percentages can be evaluated using box plot diagrams. We start this evaluation and observe high potential, however we do not test enough instances to conclude with what percentages the computation time would increase or decrease for the various characteristics.

Another interesting research topic is to evaluate the difference in objective values for problems with (limited) buffer capacity and those without the buffer capacity. This research provides insight, whether investing in (additional) buffer capacity can be profitable or not.

We can recommend many more research topics, as the zero buffer capacity is a relatively unknown issue in literature for both the HFS and FJS, as we conclude in Chapter 3. Such research opens a new topic and perspectives, therefore almost all related research and experiments enhance the literature.

6.4 Recommendations on How To Implement Our MILP

We recommend implementing the MILP in different interactive phases:

- **Phase 1** is to receive Brynild's approval for the developed scheduling method. We need the approval, as a confirmation that Brynild acknowledges the value of the new scheduling method when testing and implementing;
- **Phase 2** is to integrate the MILP in the software such that the MILP can obtain data from Brynild's ERP system;
- **Phase 3** is an iterative phase with Phase 4. In Phase 3 we conduct shadow runs; we run the MILP next to the current scheduling method. We evaluate the outcome in relation to the current method of scheduling and to signal possible discrepancies of the MILP;
- **Phase 4** is to evaluate the discrepancies and to adjust the MILP to Brynild's requirements and preferences. After adjusting the model, we conduct new shadow runs in Phase 3. When Brynild is satisfied with the outcome of the shadow runs of the MILP, we start Phase 5;
- **Phase 5** is when the MILP is fully integrated within the ERP system. Next the scheduler and the line operators need 'training' on how to work with the new scheduling method. The layout and precision of our scheduling method is different from what the employees at Brynild are used to. For the scheduler the focus is on how to use the method and how to perceive the output, the line operators need to learn how to use this schedule in practice. A human eye remains

needed to check if the theoretical model can indeed be applied in practice. However, checking the schedule is less time consuming than making the schedule from scratch. We recommend to keep the current scheduling method as a shadow run in the beginning of the implementation. The shadow run functions as a backup mechanism, and we can also use it as performance comparison. When Brynild is completely satisfied with the MILPs' functioning, the shadow runs can phase out.

We present our experiment findings to Brynild and they are excited about the results. Therefore, we already accomplish Phase 1 by gaining Brynild's approval. We jump over Phase 2 as this phase is very time consuming. We can jump over Phase 2 since in Phase 3 we enter the input data manually. We cannot yet test the integration and solely focus on the results of the schedule. We already start Phase 3 by running the MILP for Week 7 and Week 8 of 2021. We receive the feedback that the schedule should not start with Dent products, since it requires material preparation, which takes 3 hours. Therefore, in Phase 4 we add an additional constraint that states that jobs containing Dent products can only start after 3 hours in the model. Back in Phase 3, we run the model once again. This type of interactive and iterative shadow scheduling should continue until Brynild is completely satisfied with the schedules our MILP creates.

In the meanwhile, another student, Joey Klein Koerkamp, creates a new planning method for Brynild. This planning indicates what and how much to produce in what week, while considering the costs and the buffer space between the production lines. If Brynild wants to implement both the scheduling method and the new planning method, we recommend to start with the implementation of the scheduling method. Once the scheduling method is fully functional, we can also implement the planning method. There are two reasons for implementing the scheduling first. If we implement the planning first, the shadow runs for the planning are much more time consuming. We need to check the shadow runs of the planning method on feasibility, and whether there is additional space left in the schedule when the planning indicates a full week. To be able to obtain these answers, we do need to schedule the planning, which is currently done manually. The other reason is that if we start with implanting the planning method, the shadow runs and adjustments of the planning are based on the manual scheduling method. Therefore, when we implement the scheduling method, the planning has to be adjusted based on our MILP, since our method creates more optimal schedules and delivers a more accurate feasibility check.

If Brynild implements the planning method first, we require much more time. Firstly, to adjust the planning method to the manual scheduling, which itself costs more time. Moreover, once the scheduling method is implemented, we must adjust the planning method for a second time. Therefore, we recommend starting with the implementation of the scheduling method.

BIBLIOGRAPHY

(2020). Retrieved from SINTEF: https://www.sintef.no/en/

Adams, J., Balas, E., & Zawack, D. (1988). *The shifting bottleneck procedure for job shop scheduling.* Pittsburgh: Graduate School of Industrial Administration, Carnegie-Mellon University.

Aqil, S., & Allali, K. (2020). Local search metaheuristic for solving hybrid flow shop problem in slabs and beams manufacturing. *Expert Systems with Applications*.

Azzouz, A., Ennigrou, M., & Said, L. B. (2016). Flexible Job-shop Scheduling Problem with Sequence-dependent Setup Times using Genetic Algorithm. *Proceedings of the 18th International Conference of Enterprice Information Systems*, 47-53.

Bagheri, A., & Zandieh, M. (2011). Bi-criteria flexible job-shop scheduling with sequence-dependent setup times- Variable neighborhood search approach. *Journal of Manufacturing Systems*, 8-15.

Bozorgirad, M. A., & Logendran, R. (2013). *Hybrid flowshop scheduling problem with a bi-criteria objective and group technology assumption.*

Bozorgirad, M. A., & Logendran, R. (2016). A comparison of local search algorithms with population-based algorithms in hybrid flow shop scheduling problems with realistic characteristics. *International Journal Advanced Manufacturing Technology*, 1135-1151.

Brandimarte, P. (1993). Routing and schedulingin a flexible job shop by tabu search. *Annals of Operations Research*, 157-183.

Brynild Gruppen. (2020). *Produktkatalog.*

Burkard, R., Dell'Amico, M., & Martello, S. (2009). *Assignment Problems.* Philadelphia: Society for Industrial and Applied Mathematics.

Caldeira, R. H., Gnanavelbabu, A., & Solomon, J. J. (2019). *Solving the Flexible Job Shop Scheduling Problem Using a Hybrid Artificial Bee Colony Algorithm.* Springer.

Chakrabarty, D., & Goel, G. (2010). *On the approximability of budgeted allocations an improved lower bounds for submodular welfare maximization and GAP.* Society for Industrial and Applied Mathematics.

Chaudhry, I. A., & Khan, A. A. (2016). *A research survey: review of flexible job shop scheduling techniques.* Wiley.

Chen, Q., Pan, Q., Zhang, B., Ding, J., & Li, J. (2019). Effective Hot Rolling Batch Scheduling Algorithms in Compact Strip Production. *IEEE Transactions on Automation Science and Engineering*, 1933-1951.

Chiang, T.-C., & Lin, H.-J. (2013). A simple and effective evolutionary algorithm for multiobjective flexible job shop scheduling. *Internation Journal of Production Economics*, 87-98.

Chu, W. W. (1969). Optimal File Allocation in a Multiple Computre System. *IEEE Transactions on Computers*.

Dauzère-Pérès, S., Roux, W., & Lasserre, J. (1998). Multi-resource shop scheduling with resource flexibility. *European Journal of Operational Research*, 289-305.

Dekhici, L., & Belkadi, K. (2010). Operating Theatre Scheduling Under Constraints. *Journal of Applied Sciences*, 1380-1388.

Ennigrou, M., & Ghédira, K. (2008). New local diversification techniques for flexible job shop scheduling problem with a multi-agent approach. *Auton Agent Multi-Agent Syst*, 270-287.

González, M. A., Vela, C. R., & Varela, R. (2013). *An Efficient Memetic Algorithm for the Flexible Job Shop with Setup Times.* Association for the Advancement of Artificial Intelligence.

Graves, S. C. (1981). *A Review of Production Scheduling.* Cambridge: Massachusetts Institute of Technology.

Gupta, J. N. (1988). Two-Stage, Hybrid Flowshop Scheduling Problem. *The Journal of t he Operational Research Society*, 359-364.

Heerkens, H., & Winden, A. v. (2012). *Geen Probleem: Een aanpak voor alle bedrijfskundige vragen en mysteries.* Nieuwegein, The Netherlands: Van Winden Communicatie.

Hiley, A., & Julstrom, B. A. (2006). *The Quadratic Multple Knapsack Problem and Three Heuristic Approaches to it.*

Kim, T. K. (2015). *T test as a prametric statistic.* Korean Journal of Anesthesiology.

Lee, C.-G., & Ma, Z. (2003). *The Generalized Quadratic Assignment Problem.* Toronto: Department of Mechanical and Industrial Engineering, University of Toronto.

Lee, G.-C., Hong, J. M., & Choi, S.-H. (2015). *Efficient Heuristic Algorithm for Scheduling Two-Stage Hybrid Flowshop with Sequence-Dependent Setup Times.* Hindawi.

Li, S. (1997). A hybrid two-stage flowshop with part family, batch production, major and minor set-ups. *European Journal of Operational Research*, 142-156.

Li, Y., Li, X., Gao, L., & Meng, L. (2020). An improved articial bee colony algorithm for distributed heterogeneous hybrid flowshop scheduling problem with sequence-dependent setup times. *Computers & Industrial Engineering*.

Martello, S., & Toth, P. (1990). *Knapsack Problems: Algorithms and Computer Implementations.* John Wiley & Sons.

Muth, J., & Thompson, G. (1963). *Industrial Scheduling.* Englewood Cliffs: Prentice Hall.

Pezzella, F., Morganti, G., & Ciaschetti, G. (2007). *A genetic algortihm for the Flexible Job-shop Scheduling Problem.* Elsevier.

Ross, G. T., & Soland, R. M. (1975). *A branch and bound algortihm for the generalized assignmnet problem.* North-Holland Publishing Company.

Rossi, A. (2014). Flexible job shop scheduling with sequence-dependent setup and transportation times by ant colony with reinforced pheromone relationships. *Int. J. Production Economics*, 253-267.

Ruiz, R., & Maroto, C. (2006). A genetic algorithm for hybrid flowshops with sequence dependent setup times and machine eligibility. *European Journal of Operational Research*, 169, 781-800.

Ruiz, R., & Vázquez-Rodríquez, J. A. (2009). The hybrid flow shop scheduling problem. *European Journal of Operational Research*, 1-18.

Sadrzadeh, A. (2013). *Development of Both the AIS and PSO for Solving the Flexible Job Shop Scheduling Problem.* King Fahd University of Petroleum and Minerals.

Shahvari, O., & Logendran, R. (2016). Bi-criteria batch scheduling in hybrid flow shop. *2016 Industrial and Systems Engineering Research Conference*, (pp. 150-155).

Shen, L., Dauzère-Pérès, S., & Neufeld, J. S. (2017). Solving the flexible job shop scheduling problem with sequence-dependent setup times. *European Journal of Operational Research*, 503-516.

Toshev, A. (2019). *Particle Swarm Optimization and Tabu Search Hybrid Algorithm for Flexible Job Shop Scheduling Problem - Analysis of Test Results.* Sofia: Sciendo.

Wang, X., & Tang, L. (2009). A tabu search heuristic for the hybrid flowshop scheduling with finite intermediate buffers. *ScienceDirect Computers & Operations Research* , 907-918.

Wei, Q., Wu, Y., Jiang, Y., & Cheng, T. (2019). Two-machine hybrid flowshop scheduling with identical jobs: Solution algorithms and analysis of hybrid benefits. *Journal of the operational research society*, 817-826.

Williams, R. H., Pitts, S. T., & Kamery, R. H. (2004). A factory application for models and production scheduling. In A. A. Conference, *Academy of Information and Management Sceinces: Proceedings* (pp. 105-109). New Orleans.

Xia, W., & Wu, Z. (2005). *An effective hybrid optimization approach for multi-objective flexible job-shop scheduling problems.* Elsevier.

Xu, L., Jiawei, D., & Ming, H. (2017). *Research on Hybrid Cloud Particle Swarm Optimization for Multi-objective Flexible Job Shop Scheduling Problem.* International Conference on Computer Science and Network Technology.

Yaurima, V., Burtseva, L., & Tchernykh, A. (2009). Hybrid flowshop with unrelated machines, sequence-dependent setup time, availability constraints and limited buffers. *Computers & Industrial Engineering*, 1452-1463.

Yazdani, M., Amiri, M., & Zandieh, M. (2010). Flexible job-shop scheduling with parallel variable neighborhood search algorithm. *Expert Systems with Applicaitons*, 678-687.

Zandieh, M., Ghomi, S. F., & Husseini, S. M. (2006). An immune algorithm approach to hybrid flow shops scheduling with sequence-dependent setup times. *ScienceDirect*, 111-127.

Zhang, J., Ding, G., Zou, Y., Qin, S., & Fu, J. (2017). *Review of job shop scheduling research and its new perspectives under Industry 4.0.* Springer.

Zhang, Y., Li, W.-f., Storer, R. H., & Yan, X.-p. (2013). *Mathematic modeling and heuristic algorithm of hybrid flow shop with multi-jobs families and no-buffer.*

Zhang, Y., Rong, Z., & Liu, Z.-X. (2014). The integrated scheduling problem in container terminal with dual-cycle operation. *International journal simuation model*, 335-347.

APPENDIX A: GENERAL INFORMATION BRYNILD

Appendix A consist of general information regarding Brynild's current situation. Table A.0.1 consists of the family name, article number and the product name. Column 4, gives the number of intermediates that fit on 1 tray and the 5th column contains the tool number that Brynild uses to produce the intermediate.

Table A.0.1: Product information part 1

	Art. Nr	Art.navn	Antall figurer pr brett	Dyseplate
Familie A	112270-71	Godte Gomp	800	2790
	113072	Stupedama	800	2790
	113500	Påske Fruktgele	168	3245
	112958	Jordbærfisker	168	3245
	113482	Jul Fruktgelè	168	3245
	113729	Supermixgelè	168	3245
	113515	Lakrisbåter	384	2792
Familie B	112815	Knatter	800	2790
Familie C	107270	Frutti Beans xxl	476	2791
Familie D	111916	Dent Eukalyptus	660	5020
	112189	Dent Trio	660	5020
	111918	Dent Salt lakris	660	5020
	113304	Dent Crush Medium salt	660	5020
	113319	Dent Crush Pomello Jordb.	660	5020
	113943	Dent Crush Sweet Mint	660	5020
	113927	Dent Flip Lakris Karamell	168	3245/5327
	113932	Dent Flip Cola Sitron	168	3245/5327
	113940	Dent Flip Jordbær Bringebær	168	3245/5327
	113604	HF P.P Lakris Sukkerfri	476	2791
Familie K	105994	Gelepynt Freia	476	2791
(Kaldstøp)	107935	Jellymen ekstra tykke	168	3245
	106571	Myke seigmenn	168	3245
	104452	Røde Hjerter	168	3245
	106497	Sure Skumfrosker	168	3245
	113492	Figurskum-gele jul (2 lag)	168	3245/5327
	113543	Skumgele Appelsin Påske	168	3245/5327
	113542	Skumgele Cupuacu Påske	168	3245/5327
	112873	Sure Colaflasker skumgelè	168	3245/5327
	113541	Gompegelé	168	3245/5327
	113087	HF Skumgelé Solbær	168	3245/5327
	113717	Barnetimen skum/gele	168	3245/5327
	113095	HF Skumgelé Cupuacu Sommer 2017	168	3245/5327
	113494 / 113538	Søte Rakkere	168	3245
	113561/113562	Sure Rakkere	168	3245
	113563	Myke Rakkere	168	3245/5327
Familie	111132	Sure Buttons	96	5136
	111892	Skumegg	384	2792
	113496	Sure tær	156	5134

Table A.0.2 consists of the family name, article number and the product name. Column 4, gives the warm drying time and cold drying time for the new drying cabinets, and the fifth column the warm and cold drying time for the old drying cabinets. The 6th column informs us how much kilogram of an intermediate fits in 1 drying lane.

Table A.0.2: Product information part 2

			Dynaflow og Catelli		Late skapene		
	Art. Nr	Art.navn	Tid Varme	Tid Kjøl	Tid Varme	Tid Kjøl	Ferdig tørk 1 rad
Familie A	112270-71	Godte Gomp	37	10	60	24	1651
	113072	Stupedama	38	10	48	24	2056
	113500	Påske Fruktgele	38	10	48	24	2400
	112958	Jordbærfisker	38	10	48	24	2439
	113482	Jul Fruktgelè	38	10	48	24	2400
	113729	Supermixgelè	38	10	48	24	2439
	113515	Lakrisbåter	38	10	48	24	3571
Familie B	112815	Knatter	48	10	48	24	1150
Familie C	107270	Frutti Beans xxl	48	10	70	24	2506
Familie D	111916	Dent Eukalyptus	36	10	x	x	1818
	112189	Dent Trio	36	10	x	x	1818
	111918	Dent Salt lakris	48	10	x	x	1818
	113304	Dent Crush Medium salt	48	10	x	x	1871
	113319	Dent Crush Pomello Jordb.	36	10	x	x	1871
	113943	Dent Crush Sweet Mint	36	10	x	x	1871
	113927	Dent Flip Lakris Karamell	40	10	x	x	1080
	113932	Dent Flip Cola Sitron	40	10	x	x	1080
	113940	Dent Flip Jordbær Bringebær	40	10	x	x	1080
	113604	HF P.P Lakris Sukkerfri	36	10	x	x	1311
Familie K	105994	Gelepynt Freia	0	24	0	24	2838
(Kaldstøp)	107935	Jellymen ekstra tykke	0	24	0	24	3797
	106571	Myke seigmenn	0	24	0	24	2742
	104452	Røde Hjerter	0	24	0	24	3290
	106497	Sure Skumfrosker	0	24	0	24	1791
	113492	Figurskum-gele jul (2 lag)	0	48	0	48	2441
	113543	Skumgele Appelsin Påske	0	48	0	48	2618
	113542	Skumgele Cupuacu Påske	0	48	0	48	2618
	112873	Sure Colaflasker skumgelè	0	72	0	72	2723
	113541	Gompegelé	0	48	0	48	1365
	113087	HF Skumgelé Solbær	0	48	0	48	2618
	113717	Barnetimen skum/gele	0	48	0	48	1946
	113095	HF Skumgelé Cupuacu Sommer 2017	0	48	0	48	2618
	113494 / 113538	Søte Rakkere	0	48	0	48	2063
	113561/ 113562	Sure Rakkere	0	48	0	48	2257
	113563	Myke Rakkere	0	48	0	48	2082
Familie	111132	Sure Buttons	48	10	48	12	2239
	111892	Skumegg	24	10	24	10	1132
	113496	Sure tær	48	10	48	10	2540

APPENDIX B: SOLUTION APPROACHES FLEXIBLE JOB SHOPS

As we do not find not a lot of HFS problems that have the same similarities with the Brynild case we also review the FJSP with sequence-dependent set-up times (FJSP-SDST). FJSP-SDST problems have in common with the Byrnild case that the machines are non-identical and that they consider sequence dependent set-up times. However, also in this area we find no cases with zero intermediate buffer capacity. Often some sort of swarm intelligence is used like Particle Swarm by (Toshev, 2019), Ant colony by (Rossi, 2014), Hybrid Cloud Particle Swarm by (Xu, Jiawei, & Ming, 2017), and Bee colony by (Caldeira, Gnanavelbabu, & Solomon, 2019). These swarm intelligence applications are most of the time integrated with a GA or TS, therefore we focus on the GA and TS. The three most used and evaluated methods to solve the FJSP-SDST are Tabu Search, Genetic Algorithms and Variable Neighbourhood Search.

A FJSP-SDST formulated by Azzouz, Ennigrou & Said (2016) with the objective to minimize make span is solved by a Genetic Algorithm (GA). The GA is evaluated for two types of objectives, make span and Aggregate Objective Function. In this paper they compare GA to a VNS, which is proposed by Bagheri & Zandieh (2011), an adapted TS proposed by Ennigrou & Ghédira (2008), Artificial Immune System and Particle Swarm Optimization from Sadrzadeh (2013). The relative percentage deviation is used to compare the performance of the algorithms. GA performs best in most instances, however in 3 out of 20 instances VNS performs the best and for 1 instance TS performs the best.

Shen, Dauzère-Pérès & Neufeld (2017) solve a FJSP-SDST with a TS and the objective is to minimize the make span. Four moves are identified, 2 focus on the resequencing on the same machine and the other 2 focus on reassignments to another machine. González, Vela & Varela (2013) propose a neighbourhood, which is embedded into a GA hybridized with TS for a FJSP-SDST. They find that a TS combined with a GA works very well. For the smaller instances of 10 jobs the TS by itself finds the same values or very close values to the GA+TS method.

Bagheri & Zandieh (2011) consider a FJSP with sequence-dependent set-up times to minimize the make span and mean tardiness. To solve this FJSP problem a variable neighbourhood search algorithm is proposed. First the operations are assigned to a machine. After this assignment, the sequence is determined. Multiple random feasible solutions are produced and the best one is chosen. The method is evaluated against the parallel variable neighbourhood search algorithm of Yazdani, Amiri & Zandieh (2010) and the genetic algorithm of Pezzella, Morganti & Ciaschetti (2007). The relative percentage deviation is used as an evaluation. the variable neighbourhood search performs better than the adapted parallel variable neighbourhood search and genetic algorithm for all instances.

Appendix C: Decisions concerning the model

We explain in the appendix the 2 important decisions we make regarding the main model. Firstly, the pre-processing of the input data for the model. We discuss the various options we consider and the advantages and disadvantages of the options. Secondly, the objective of the model. We discuss different objectives and the advantages and disadvantages of the objectives.

Pre-processing input data

Data pre-processing options

In this section we determine how detailed we want our mathematical model to be. The degree of the details influence the number of constraints, variables, and to what extent we need to pre-process the input data. The raw input data for the Støperi 1 scheduling problem is the output data of the week planning. This planning indicates the quantity in kg of every intermediate to be produced each week. In this research we consider using 3 different options on how to use and pre-process the input data.

1. The input data is the direct output data of the planning;
2. We pre-process the input data slightly; we divide intermediates that are too large, so 1 intermediate always fit in the drying cabinet;
3. We pre-process the input data fully; we divide intermediates that are too large, and combine intermediates that can dry together, thus 1 or multiple intermediates fit in the drying cabinet together.

We examine different input data possibilities, as the input data has a large effect on our mathematical model. The 3 options have in common that the following conditions should be taken into account:

- No buffer points;
- Sequence dependent set-up times;
- Non-identical parallel machines;
- Production time windows.

There are zero buffer points in the model, every process passes through in a continuous flow. The sequence dependent set-up times depends on whether the intermediate is the same or if the intermediates make use of the same tool. With the non-identical parallel machines, the intermediates have different throughput times for the drying cabinets. Last, we require the production time windows to indicate the 2-shift schedules and to end the 3-shift time schedules.

The 3 different options for data input, have different levels of flexibility and scheduling difficulties, which we discuss in the subsections below.

Input data is the planning output data

In this first option, the input data has the same format as the output data from the planning; the different intermediates with their quantity for every week. When the input of the model is the output of the planning the following additions must be taken into account:

- Batch sizing;
- Intermediates drying together.

It can occur that the quantity of the planned intermediate is too large and does not fit in 1 drying cabinet. Therefore, the mathematical model needs to be able to make batches for every intermediate. The batches have a maximum quantity in kg that fit in the drying cabinets. The quantity depends on each intermediate and each drying cabinet.

Next to batch sizing, this input data also requires the option in which intermediates could dry in the same drying cabinet. The mathematical model has to determine whether intermediates dry together, if the intermediates fit together and can dry together in the drying cabinets and recalculate the utilization time of the drying cabinet. The drying end time for both intermediates becomes the same and the drying only starts when the second intermediate enters the drying cabinet.

A disadvantage of having the output data from the planning as input data is that the model has a large and complex solution space. We need various constraints and decision variables for this model to work. Since the batch size still requires determination and is not integer, continuum of possibilities arise. This mathematical model requires a great time to compute, at least multiple days, if the instances are even small enough to be solvable. Another disadvantage is that the output could be hard to apply in practice. It is difficult to produce the exact quantity from the schedule, it often occurs that Brynild produces more, or less than they intend.

Although using the output data directly as input imposes disadvantages, there are also two important advantages. The first advantage is that the data does not need pre-processing. Another advantage is that the model creates the most optimal, given the planning as input, schedule since we do not reduce the flexibility of the model in any way.

Slight pre-processing of the input data: 1 intermediate fits in the drying cabinet
In the second option for the input data, we alter the intermediates quantity such that we make batches before scheduling that always fit in every drying cabinet. There are multiple ways to divide the production quantities. One is to evenly distribute the quantity over the number of batches we minimal need. Another option could be to completely 'fill' multiple batches and have 1 batch left containing the rest. We must research what would generally work the best for Brynild's scheduling problem. We still need to consider the following addition when the input of the model are pre-processed batches containing 1 intermediate that fits in every drying cabinet:

- Intermediates drying together.

The advantages of determining the batch sizes before scheduling is that we sperate the intermediates beforehand, so the final output is more useful in practice. This is the case since the scheduler is able to do an additional check on feasibility, before using the input data in the model. The second advantage is that we need less constraints and decision variables, since the model no longer needs to separate the input data in batches to fit the drying cabinets. Therefore, the model requires less computation.

A disadvantage of pre-determining the batches is the loss of flexibility. By pre-determining the batches, it could be that the model can no longer create the most optimal schedule. Another disadvantage is that the mathematical model still needs to have constraints and variables to dry different intermediates together. Therefore, the parallel machines must be able to process multiple batches at the same time. One of the main assumptions of a hybrid flow shop we mention in Chapter 3 is that each machine can only process 1 operation at the time and each job can be processed by only 1 machine at the time.

Pre-processing of the input data: 1 or multiple intermediates fit in the drying cabinet together
For the third option we further alter the output data of the planning than in Option 2. As opposed to Option 2, batch sizes made in Option 3 consider the additional possibility of batches drying together. This would allow for consecutive production of the batch combinations, such that we minimize the waiting time in an open drying cabinet. This additional alteration to produce the intermediates consecutively when they dry together is not farfetched, as the first intermediate must always wait for the other intermediates before the drying cabinet can begin. Also, in practice it is far from ideal to have liquid intermediates waiting for a long time in an open drying cabinet, since the quality of the

product reduces. If we alter the input data for Option 3, no additions need to be taken into account. However, we do need to consider the 2 Catelli drying cabinets as one and the 2 Dynaflow cabinets as one. We consider 5 drying cabinets with the same capacity instead of 7 drying cabinets with varying capacities. Because if we consider cabinets with varying capacities, we cannot know how large a job can be. Therefore, all jobs need to have the size of the smallest cabinet and the model becomes far from optimal, or should still consider jobs drying together.

The main advantage of pre-processing the input data by making batches containing intermediates to dry together is that the model requires less constraints and decision variables. The constraints in the model no longer need to consider, batches drying together in a drying cabinet since we already combine the intermediates in 1 batch. Therefore, this model can compute faster than the models of the other two options. Another advantage is that the undesired outcomes are no longer in the solution space of the model. With undesired outcomes we mean outcomes undesired in practice, for example where the first intermediate must wait over night before the second intermediate starts producing, since this is not good for the quality of the products.

The disadvantage is more loss of flexibility than with Option 2. And that we consider 5 drying cabinets equal in capacity instead of 7 uneven in capacity. Thus, there is even a higher probability that, given the planning as input, the theoretical optimal solution is not in the solution space of the model when using the input data of Option 3.

Determination of the input data

In Section 4.1.2 we examine 3 options for pre-processing of the input data. Option 1 needs the most variables and constraints as this option considers the least pre-processing. Option 3 requires the most pre-processing and therefore needs the least variables and constraints. Table C.0.1, shows what we need to consider for each data input option.

Table C.0.1: Input data processing option 1, 2, and 3 summarized

	Option 1	**Option 2**	**Option 3**
No buffer points	X	X	X
Sequence dependent set-up times	X	X	X
Non-identical parallel machines	X	X	X
Production time windows	X	X	X
Batch sizing	X	X	
Intermediates drying together	X		

After considering the advantages and disadvantages of the 3 options for the input data, we chose to use Option 3: we batch 1 or multiple intermediates together to fit in the drying cabinet. In Option 3 the number of constraints, variables and the solution space is more manageable. Due to these simplifications for Option 3, see Table C.0.1, the probability of being able to solve the mathematical problem in the least amount of time is the highest. The loss of flexibility is unfortunate for reaching the optimality in the theoretical schedule. However, this loss of flexibility makes the theoretical schedule of Option 3 most feasible in practice.

OBJECTIVE VALUE

Given the constraints from Chapter 4, we can optimize the schedule of Brynild using various objectives. Brynild indicates that they focus on minimizing the individual throughput time. Their reasoning for this that with lower throughput times they are able to produce more products. However, minimizing individual throughput time does not always result in being able to produce more.

We examine 3 objectives, of which the drying time of the job is central. We look specifically at the end time of jobs and not the throughput time. Otherwise, in theory the schedule can obtain the same objective starting on Monday or Tuesday as long as the throughput is the same. This is not desirable in practice; therefore, we disregard objectives based on throughput times. More specifically we analyse 3 objectives that all focus on the total make span, thus end-drying times of the jobs.

Objective 1 is the minimization of the make span of all jobs:

$$\min\left[\max_{j,c} D_{j,c}\right]$$

We minimize the latest time a job finishes drying. This objective does not minimize every individual throughput time of each job. The objective spreads the work force as evenly as possible for the drying cabinets to finish with all jobs as early in the week as possible.

Objective 2 is to minimize the summation over the make span of each job:

$$\min \sum_{j \in J} D_j$$

This objective minimizes the make span for each job. Therefore, this objective results in scheduling the shorter jobs first and the longer jobs at the end of the week.

Objective 3 is to minimize the summation over the latest end times at all cabinets:

$$\min\left[\sum_{c \in C} \max D_{j,c}\right]$$

This objective results in minimizing the throughput time per job the most of the 3 objectives. Because the jobs that have different drying times on the drying cabinets we schedule on the cabinets with the shortest time.

These 3 objectives have advantages and disadvantages in comparison to each other.

As a consequence of Objective 2, that schedules shorter jobs first and longer jobs at the end, the solution space becomes limited and does not necessarily results in a schedule that Brynild desires. Producing the short jobs first and the longer jobs later does not result in producing more than producing the long jobs first and the shorter jobs later. Therefore, we do not select Objective 2.

Objective 3 does not differentiate between; 2 drying cabinets both finishing on Thursday or 1 on Wednesday and 1 on Friday. While in objective 1 the drying both finishes on Thursday. Both outcomes have pros and cons. The second option, drying finishing on Wednesday and on Friday, could be more desirable when the model runs just once and the change in jobs or unplanned jobs are added manually. Because the drying cabinets cannot start emptying at the same time, therefore it would be better when they end at different times. However, the drawback of Objective 3 is that this objective fills the new drying cabinets the most as these give the least throughput time. In terms of flexibility later in the week, we do not desire this.

The reason that we do not select Objective 3 is that Brynild indicates that schedules may change due to emergency jobs or other unforeseen complications. In this case it would be better to do as much as possible in the beginning of the week. Then, when changes occur, Brynild should run the model again. Therefore, it wouldn't matter that 2 drying cabinets finish at the same time. As the model runs again to determine a new schedule, or the production week finish early.

So, we select Objective 1 as the objective for our MILP. Objective 1, does not directly minimize the throughput time for the jobs individually, however Objective 1 does minimize the total make span of 1 week. Next, with a schedule based on Objective 1, it is easier to anticipate on unforeseen changes and makes the schedule more flexible during the week.

APPENDIX D: MATHEMATICAL MODELS

We construct Appendix D according to Table D.0.1. As addition we add Model 9, which consists of 2 separate models, one for the assignment of jobs to parallel machines and one for the sequencing of jobs.

Table D.0.1: Different MILPs without buffer capacity

Model	**Non-identical parallel machines**	**Sequence dependent set-up times**	**Production time windows**
1	X	X	X
2	X	X	
3	X		
4		X	
5			X
6		X	X
7	X		X
8			

MODEL 1: MAIN MILP

Indices

$j, i \in J$	Jobs, including 2 dummy jobs, dummy job 1 at the beginning of the schedule and dummy job 2 at the end of each schedule
$c \in C$	Drying cabinets
$p \in P$	Processing days

Parameters

m_j	Moulding time of job j in hours
$v_{j,c}$	Drying time of job j in drying cabinet c in hours
$s_{i,j}$	Set-up time of job j after job i for Stage 1 in hours
k_j	Cooking time for job j in hours
o_p	First possible non-negative start time of processing day p (day 1 has as start time 0)
l_p	Latest possible non-negative end time in hour of processing day p
$g_{j,c}$	Binary parameter, indicating whether job j, can be dried in drying cabinet c (g=1), or not (g=0)
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
T_j	Duration of job j in Stage 1 (set-up, cooking and moulding)
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$W_{j,p}$	Indicates whether job j starts processing at day p for $W_{j,p}$=1, otherwise 0
$X_{i,j}$	Indicates whether job j is sequenced somewhere after job i in Stage 1 and therefore in Stage 2 for $X_{i,j}$=1, otherwise 0
$Y_{j,c}$	Indicates whether job j is assigned to cabinet c for $Y_{j,c}$=1, otherwise 0
$Z_{i,j}$	Indicates whether job j is sequenced directly after job i in Stage 1 for $Z_{i,j}$, otherwise 0

Constraints

$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$

$X_{j,j} = 0 \qquad \forall i = j$

$Z_{i,j} \leq X_{i,j} \qquad \forall i, j \quad i \neq j$

$\sum_{i \in J} Z_{i,j} = 1 \qquad \forall j \neq dummy\, 1 \quad i \neq j$

$\sum_{j \in J} Z_{i,j} = 1 \qquad \forall i \neq dummy\, 2 \quad i \neq j$

$B_j = 0 \qquad j = dummy\, 1$

$\sum_{j \in J} X_{i,j} = 0 \qquad i = dummy\, 2$

$T_j = k_j + m_j + \sum_{i \in J}(s_{i,j} * Z_{i,j}) \qquad \forall j$

$E_j = B_j + T_j$ $\forall j$

$E_j - E_i + M(1 - Z_{i,j}) \geq T_j$ $\forall i, j \quad i \neq j$

$\sum_{p \in P} W_{j,p} = 1$ $\forall j$

$\sum_{p \in P}(W_{j,p} * o_p) \leq B_j$ $\forall j$

$E_j \leq \sum_{p \in P}(W_{j,p} * l_p)$ $\forall j$

$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c})$ $\forall i, j, c \quad i \neq j$

$N_j = E_j - m_j$ $\forall j$

$D_{j,c} = E_j + v_{j,c} * Y_{j,c}$ $\forall j, c$

$\sum_c Y_{j,c} = 1$ $\forall j \neq dummy\ 1\ and\ dummy\ 2$

$Y_{j,c} \leq g_{j,c}$ $\forall j, c$

$W_{j,p}, \quad X_{i,j}, \quad Y_{j,c}, \quad Z_{i,j} \quad \in \quad \{0,1\}$

$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c}, \quad \geq \quad 0$

MODEL 2: WITHOUT PRODUCTION TIME WINDOWS

The difference with the main MILP is that we no longer need the index p. With the disappearance of this index, all the parameters, variables and constraints that have a connection to this index disappear as well.

Indices

$j, i \in J$	Jobs, including 2 dummy jobs, dummy job 1 at the beginning of the schedule and dummy job 2 at the end of each schedule
$c \in C$	Drying cabinets

Parameters

m_j	Moulding time of job j in hours
$v_{j,c}$	Drying time of job j in drying cabinet c in hours
$s_{i,j}$	Set-up time of job j after job i for Stage 1 in hours
k_j	Cooking time for job j in hours
$g_{j,c}$	Binary parameter, indicating whether job j, can be dried in drying cabinet c (g=1), or not (g=0)
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
T_j	Duration of job j in Stage 1 (set-up, cooking and moulding)
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$X_{i,j}$	Indicates whether job j is sequenced somewhere after job i in Stage 1 and therefore in Stage 2 for $X_{i,j}$=1, otherwise 0
$Y_{j,c}$	Indicates whether job j is assigned to cabinet c for $Y_{j,c}$=1, otherwise 0
$Z_{i,j}$	Indicates whether job j is sequenced directly after job i in Stage 1 for $Z_{i,j}$, otherwise 0

Constraints

$X_{i,j} + X_{j,i} = 1$	$\forall i, j \quad i \neq j$
$X_{j,j} = 0$	$\forall i = j$
$Z_{i,j} \leq X_{i,j}$	$\forall i, j \quad i \neq j$
$\sum_{i \in J} Z_{i,j} = 1$	$\forall j \neq dummy\ 1 \quad i \neq j$
$\sum_{j \in J} Z_{i,j} = 1$	$\forall i \neq dummy\ 2 \quad i \neq j$
$B_j = 0$	$j = dummy\ 1$
$\sum_{j \in J} X_{i,j} = 0$	$i = dummy\ 2$
$T_j = k_j + m_j + \sum_{i \in J}(s_{i,j} * Z_{i,j})$	$\forall j$

$E_j = B_j + T_j$ $\quad\quad$ $\forall j$

$E_j - E_i + M(1 - Z_{i,j}) \geq T_j$ $\quad\quad$ $\forall i, j \quad i \neq j$

$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c})$ $\quad\quad$ $\forall i, j, c \quad i \neq j$

$N_j = E_j - m_j$ $\quad\quad$ $\forall j$

$D_{j,c} = E_j + v_{j,c} * Y_{j,c}$ $\quad\quad$ $\forall j, c$

$\sum_c Y_{j,c} = 1$ $\quad\quad$ $\forall j \neq dummy\ 1\ and\ dummy\ 2$

$Y_{j,c} \leq g_{j,c}$ $\quad\quad$ $\forall j, c$

$X_{i,j}, \quad Y_{j,c}, \quad Z_{i,j}$ $\quad\quad$ $\in \quad \{0,1\}$

$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c},$ $\quad\quad$ $\geq \quad 0$

Model 3: Without SDST and No Production Time Windows

The difference in comparison to the main MILP are the same as for Model 2, the disappearance of index p and that we do not need the variable for the exact sequence. Therefore, we no longer need the dummy jobs. We do not need this variable as there are no longer any sequence dependent set up times. We give every job a standard set-up of the maximum set-up time (two hours).

Indices

$j, i \in J$	Jobs
$c \in C$	Drying cabinets

Parameters

m_j	Moulding time of job j in hours
$v_{j,c}$	Drying time of job j in drying cabinet c in hours
s	Set-up time is 2
k_j	Cooking time for job j in hours
t_j	Duration of job j in Stage 1; $t_j = k_j + m_j + s$
$g_{j,c}$	Binary parameter, indicating whether job j, can be dried in drying cabinet c (g=1), or not (g=0)
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$X_{i,j}$	Binary variable indicating whether job j is sequenced somewhere after job i
$Y_{j,c}$	Binary variable indicating whether job j is assigned to cabinet c

Constraints

$$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$$

$$X_{j,j} = 0 \qquad \forall j$$

$$E_j = B_j + t_j \qquad \forall j$$

$$E_j - E_i + M(1 - X_{i,j}) \geq t_j \qquad \forall i, j \quad i \neq j$$

$$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c}) \qquad \forall i, j, c \quad i \neq j$$

$$N_j = E_j - m_j \qquad \forall j$$

$$D_{j,c} = E_j + v_{j,c} * Y_{j,c} \qquad \forall j, c$$

$$\textstyle\sum_c Y_{j,c} = 1 \qquad \forall j$$

$$Y_{j,c} \leq g_{j,c} \qquad \forall j, c$$

$X_{i,j}, \quad Y_{j,c} \qquad \in \qquad \{0,1\}$

$B_j, \quad N_j, \quad E_j, \quad D_{j,c}, \qquad \geq \qquad 0$

MODEL 4: WITH IDENTICAL PARALLEL MACHINES AND NO PRODUCTION TIME WINDOWS

The difference with the main MILP is that we no longer need the index p. With the disappearance of this index, all the parameters, variables and constraints that have a connection to this index disappear as well. Next, we do not need index c for the drying time as all drying times for every cabinet are the same is solely dependent on the job and we also do not need the binary parameter to determine whether a job can dry in a specific drying cabinet, since the drying cabinets are identical.

Indices

$j, i \in J$	Jobs, including 2 dummy jobs, dummy job 1 at the beginning of the schedule and dummy job 2 at the end of each schedule
$c \in C$	Drying cabinets

Parameters

m_j	Moulding time of job j in hours
v_j	Drying time of job j
$s_{i,j}$	Set-up time of job j after job i for Stage 1 in hours
k_j	Cooking time for job j in hours
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
T_j	Duration of job j in Stage 1 (set-up, cooking and moulding)
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$X_{i,j}$	Indicates whether job j is sequenced somewhere after job i in Stage 1 and therefore in Stage 2 for $X_{i,j}$=1, otherwise 0
$Y_{j,c}$	Indicates whether job j is assigned to cabinet c for $Y_{j,c}$=1, otherwise 0
$Z_{i,j}$	Indicates whether job j is sequenced directly after job i in Stage 1 for $Z_{i,j}$, otherwise 0

Constraints

$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \neq j$

$X_{j,j} = 0 \qquad \forall i = j$

$Z_{i,j} \leq X_{i,j} \qquad \forall i, j \quad i \neq j$

$\sum_{i \in J} Z_{i,j} = 1 \qquad \forall j \neq dummy\ 1 \quad i \neq j$

$\sum_{j \in J} Z_{i,j} = 1 \qquad \forall i \neq dummy\ 2 \quad i \neq j$

$B_j = 0 \qquad j = dummy\ 1$

$\sum_{j \in J} X_{i,j} = 0 \qquad i = dummy\ 2$

$T_j = k_j + m_j + \sum_{i \in J} (s_{i,j} * Z_{i,j}) \qquad \forall j$

$E_j = B_j + T_j$ $\forall j$

$E_j - E_i + M(1 - Z_{i,j}) \geq T_j$ $\forall i, j \quad i \neq j$

$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c})$ $\forall i, j, c \quad i \neq j$

$N_j = E_j - m_j$ $\forall j$

$D_{j,c} = E_j + v_j * Y_{j,c}$ $\forall j, c$

$\sum_c Y_{j,c} = 1$ $\forall j \neq dummy\, 1\ and\ dummy\, 2$

$X_{i,j}, \quad Y_{j,c}, \quad Z_{i,j}$ $\in \quad \{0,1\}$

$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c},$ $\geq \quad 0$

MODEL 5: WITH IDENTICAL PARALLEL MACHINES AND NO SDST

In comparison to the main MILP, we do not need index c for the drying time as all drying times for every cabinet are the same is solely dependent on the job and we also do not need the binary parameter to determine whether a job can dry in a specific drying cabinet, since the drying cabinets are identical. Next, we do not need the variable for the exact sequence. Therefore, the dummy jobs are not needed any longer. We do not need this variable as there are no longer any sequence dependent set up times. We give every job a standard set-up of the maximum set-up time (two hours).

Indices

$j, i \in J$	Jobs
$c \in C$	Drying cabinets
$p \in P$	Processing days

Parameters

m_j	Moulding time of job j in hours
v_j	Drying time of job j
s	Set-up time is 2
k_j	Cooking time for job j in hours
t_j	Duration of job j in Stage 1; $t_j = k_j + m_j + s$
o_p	First possible start time of processing day p (day 1 has as start time 0)
l_p	Latest possible time in hour of processing day p
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$W_{j,p}$	Binary variable indicating whether job j starts processing at day p
$X_{i,j}$	Binary variable indicating whether job j is sequenced somewhere after job i
$Y_{j,c}$	Binary variable indicating whether job j is assigned to cabinet c

Constraints

$$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$$

$$X_{j,j} = 0 \qquad \forall j$$

$$E_j = B_j + t_j \qquad \forall j$$

$$E_j - E_i + M(1 - X_{i,j}) \geq t_j \qquad \forall i, j \quad i \neq j$$

$$\sum_p W_{j,p} = 1 \qquad \forall j$$

$$\sum_p (W_{j,p} * o_p) \leq B_j \qquad \forall j$$

$$E_j \leq \sum_p (W_{j,p} * l_p) \qquad \forall j$$

$D_{i,c} \leq N_j + M\left(2 - X_{i,j} - Y_{j,c}\right)$	$\forall i, j, c \quad i \neq j$
$N_j = E_j - m_j$	$\forall j$
$D_{j,c} \ = E_j + v_j * Y_{j,c}$	$\forall j, c$
$\sum_c Y_{j,c} = 1$	$\forall j$
$W_{j,p}, \quad X_{i,j}, \quad Y_{j,c}$	$\in \qquad \{0,1\}$
$B_j, \quad N_j, \quad E_j, \quad D_{j,c},$	$\geq \qquad 0$

MODEL 6: WITH IDENTICAL PARALLEL MACHINES, SDST, AND PRODUCTION TIME WINDOWS

In comparison to the main MILP, we do not need index c for the drying time as all drying times for every cabinet are the same is solely dependent on the job and we also do not need the binary parameter to determine whether a job can dry in a specific drying cabinet, as the drying cabinets are identical.

Indices

$j, i \in J$	Jobs, including 2 dummy jobs, dummy job 1 at the beginning of the schedule and dummy job 2 at the end of each schedule
$c \in C$	Drying cabinets
$p \in P$	Processing days

Parameters

m_j	Moulding time of job j in hours
v_j	Drying time of job j
$s_{i,j}$	Set-up time of job j after job i for Stage 1 in hours
k_j	Cooking time for job j in hours
o_p	First possible non-negative start time of processing day p (day 1 has as start time 0)
l_p	Latest possible non-negative end time in hour of processing day p
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
T_j	Duration of job j in Stage 1 (set-up, cooking and moulding)
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$W_{j,p}$	Indicates whether job j starts processing at day p for $W_{j,p}$=1, otherwise 0
$X_{i,j}$	Indicates whether job j is sequenced somewhere after job i in Stage 1 and therefore in Stage 2 for $X_{i,j}$=1, otherwise 0
$Y_{j,c}$	Indicates whether job j is assigned to cabinet c for $Y_{j,c}$=1, otherwise 0
$Z_{i,j}$	Indicates whether job j is sequenced directly after job i in Stage 1 for $Z_{i,j}$, otherwise 0

Constraints

$$X_{i,j} + X_{j,i} = 1 \qquad \forall i,j \quad i \neq j$$

$$X_{j,j} = 0 \qquad \forall i = j$$

$$Z_{i,j} \leq X_{i,j} \qquad \forall i,j \quad i \neq j$$

$$\sum_{i \in J} Z_{i,j} = 1 \qquad \forall j \neq dummy\, 1 \quad i \neq j$$

$$\sum_{j \in J} Z_{i,j} = 1 \qquad \forall i \neq dummy\, 2 \quad i \neq j$$

$$B_j = 0 \qquad j = dummy\, 1$$

$$\sum_{j \in J} X_{i,j} = 0 \qquad i = dummy\, 2$$

$T_j = k_j + m_j + \sum_{i \in J}(s_{i,j} * Z_{i,j})$ $\forall j$

$E_j = B_j + T_j$ $\forall j$

$E_j - E_i + M(1 - Z_{i,j}) \geq T_j$ $\forall i, j \quad i \neq j$

$\sum_{p \in P} W_{j,p} = 1$ $\forall j$

$\sum_{p \in P}(W_{j,p} * o_p) \leq B_j$ $\forall j$

$E_j \leq \sum_{p \in P}(W_{j,p} * l_p)$ $\forall j$

$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c})$ $\forall i, j, c \quad i \neq j$

$N_j = E_j - m_j$ $\forall j$

$D_{j,c} = E_j + v_j * Y_{j,c}$ $\forall j, c$

$\sum_c Y_{j,c} = 1$ $\forall j \neq dummy\ 1\ and\ dummy\ 2$

$W_{j,p}, \quad X_{i,j}, \quad Y_{j,c}, \quad Z_{i,j}$ $\in \quad \{0,1\}$

$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c},$ $\geq \quad 0$

MODEL 7: WITHOUT SEQUENCE DEPENDENT SET-UP TIMES

The difference with the main MILP is that we do not need the variable for the exact sequence and therefore we do not need the dummy jobs any longer. We do not need this variable as there are no longer any sequence dependent set up times. We give every job a standard set-up of the maximum set-up time (two hours).

Indices

$j, i \in J$	Jobs
$c \in C$	Drying cabinets
$p \in P$	Processing days

Parameters

m_j	Moulding time of job j in hours
$v_{j,c}$	Drying time of job j in drying cabinet c in hours
s	Set-up time is 2
k_j	Cooking time for job j in hours
t_j	Duration of job j in Stage 1; $t_j = k_j + m_j + s$
o_p	First possible start time of processing day p (day 1 has as start time 0)
l_p	Latest possible time in hour of processing day p
$g_{j,c}$	Binary parameter, indicating whether job j, can be dried in drying cabinet c (g=1), or not (g=0)
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$W_{j,p}$	Binary variable indicating whether job j starts processing at day p
$X_{i,j}$	Binary variable indicating whether job j is sequenced somewhere after job i
$Y_{j,c}$	Binary variable indicating whether job j is assigned to cabinet c

Constraints

$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$

$X_{j,j} = 0 \qquad \forall j$

$E_j = B_j + t_j \qquad \forall j$

$E_j - E_i + M(1 - X_{i,j}) \geq t_j \qquad \forall i, j \quad i \neq j$

$\sum_p W_{j,p} = 1 \qquad \forall j$

$\sum_p (W_{j,p} * o_p) \leq B_j \qquad \forall j$

$E_j \leq \sum_p (W_{j,p} * l_p) \qquad \forall j$

$$D_{i,c} \leq N_j + M\left(2 - X_{i,j} - Y_{j,c}\right) \qquad \forall i,j,c \quad i \neq j$$

$$N_j = E_j - m_j \qquad \forall j$$

$$D_{j,c} = E_j + v_{j,c} * Y_{j,c} \qquad \forall j,c$$

$$\sum_c Y_{j,c} = 1 \qquad \forall j$$

$$Y_{j,c} \leq g_{j,c} \qquad \forall j,c$$

$$W_{j,p}, \quad X_{i,j}, \quad Y_{j,c} \qquad \in \qquad \{0,1\}$$

$$B_j, \quad N_j, \quad E_j, \quad D_{j,c}, \qquad \geq \qquad 0$$

Model 8: Basis Model for No Buffer Capacity

Indices

$j, i \in J$	Jobs
$c \in C$	Drying cabinets

Parameters

m_j	Moulding time of job j in hours
v_j	Drying time of job j
s	Set-up time is 2
k_j	Cooking time for job j in hours
t_j	Duration of job j in Stage 1; $t_j = k_j + m_j + s$
M	Very large number

Variables

Non-binary variables

B_j	Begin time of job j
N_j	Start moulding time of job j
E_j	End time of job j in Stage 1, after moulding
$D_{j,c}$	End time of job j in cabinet c at Stage 2, after drying

Binary variables

$X_{i,j}$	Binary variable indicating whether job j is sequenced somewhere after job i
$Y_{j,c}$	Binary variable indicating whether job j is assigned to cabinet c

Constraints

$$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$$

$$X_{j,j} = 0 \qquad \forall j$$

$$E_j = B_j + t_j \qquad \forall j$$

$$E_j - E_i + M\left(1 - X_{i,j}\right) \geq t_j \qquad \forall i, j \quad i \neq j$$

$$D_{i,c} \leq N_j + M\left(2 - X_{i,j} - Y_{j,c}\right) \qquad \forall i, j, c \quad i \neq j$$

$$N_j = E_j - m_j \qquad \forall j$$

$$D_{j,c} = E_j + v_j * Y_{j,c} \qquad \forall j, c$$

$$\sum_c Y_{j,c} = 1 \qquad \forall j$$

$$X_{i,j}, \quad Y_{j,c} \qquad \in \quad \{0,1\}$$

$$B_j, \quad N_j, \quad E_j, \quad D_{j,c}, \qquad \geq \quad 0$$

MODEL 9: ASSIGNMENT & SEQUENCING SEPARATED

This model consists of 2-separate models namely the assignment of jobs to the drying cabinets. We use the output (the assignment) as input parameter for the second model, which determines the sequence.

Assignment

The assignment model has two constraints the same as the main model. However, now only the assignment is the outcome of this model. Therefore, we need an 'intermediate' objective.

Indices (Sets)

$j \in J$	Jobs
$c \in C$	Drying cabinets

Parameters

m_j	Moulding time of job j in hours
$v_{j,c}$	Drying time of job j in drying cabinet c in hours
$g_{j,c}$	Binary parameter, indicating whether job j, can be dried in drying cabinet c (g=1), or not (g=0)

Variables

R_c	Total time as quantity assigned to cabinet c
$Y_{j,c}$	Binary variable indicating whether job j is assigned to cabinet c

Constraints

$$\sum_c Y_{j,c} = 1 \qquad \forall j$$

$$Y_{j,c} \leq g_{j,c} \qquad \forall j, c$$

$$R_c = \sum_j (Y_{j,c} * (m_j + v_{j,c})) \qquad \forall c$$

$$Y_{j,c} \qquad \in \qquad \{0,1\}$$

$$R_c \qquad \geq \qquad 0$$

Objective

We cannot exactly know how full a drying cabinet can be. It could be said a drying cabinet can be full for maximum a week as we want the cabinets to finish before the start of the consecutive week. The problem is that when drying cabinets are full for precisely one week, at least four of the five drying cabinets do not be finish before the end of the week as these drying cabinets have to start later with drying than the drying cabinet that gets assigned the first job that is produced. With the same reasoning, we do not know beforehand that when a job finishes drying, a new job is already produced and could directly enter the drying cabinet. Thus, we do not know exactly the maximum amount of time a drying cabinet has, to obtain a feasible solution. Therefore, we decide on an objective where this problem as a very low probability to occur.

As objective we decide to minimize the maximum total time assigned to the drying cabinets:

$$\min \left[\max_c R_c \right]$$

Therefore, no cabinet is much fuller than another, which is in line with the objective of the sequencing the jobs. After, assigning the jobs to the drying cabinets to distribute the workload the sequencing of the jobs is determined by the following mathematical model.

Sequencing

Indices

$j, i \in J$ Jobs
$c \in C$ Drying cabinets
$p \in P$ Processing days

Parameters

m_j Moulding time of job j in hours
$v_{j,c}$ Drying time of job j in drying cabinet c in hours
$s_{i,j}$ Set-up time of job j after job i in hours
k_j Cooking time for job j in hours
o_p First possible start time of processing day p (day 1 has as start time 0)
l_p Latest possible time in hour of processing day p
$y_{j,c}$ Binary parameter, indicating whether job j is dried in drying cabinet c (g=1), or not (g=0)
M Very large number

Variables

B_j Begin time of job j
T_j Duration of job j in Stage 1
N_j Start moulding time of job j
E_j End time of job j in Stage 1, after moulding
$D_{j,c}$ End time of job j in cabinet c at Stage 2, after drying

$W_{j,p}$ Binary variable indicating whether job j starts processing at day p
$X_{i,j}$ Binary variable indicating whether job j is sequenced somewhere after job i
$Z_{i,j}$ Binary variable indicating whether job j is sequenced directly after job i

Constraints

A lot of constraints are the same as in the main model. The difference is that the assignment of jobs to the drying cabinets is now a parameter instead of a variable.

$X_{i,j} + X_{j,i} = 1 \qquad \forall i, j \quad i \neq j$

$X_j = 0 \qquad \forall j$

$Z_{i,j} \leq X_{i,j} \qquad \forall i, j \quad i \neq j$

$\sum_i Z_{i,j} = 1 \qquad \forall j \neq dummy\, 1$

$\sum_j Z_{i,j} = 1 \qquad \forall i \neq dummy\, 2$

$B_j = 0 \qquad j = dummy\, 1$

$\sum_{j \in J} X_{i,j} = 0 \qquad i = dummy\, 2$

$T_j = k_j + m_j + \sum_i (s_{i,j} * Z_{i,j}) \qquad \forall j$

$E_j = B_j + T_j \qquad \forall j$

$E_j - E_i + M(1 - Z_{i,j}) \geq T_j$ $\quad\quad \forall i, j \quad i \neq j$

$\sum_p W_{j,p} = 1$ $\quad\quad \forall j$

$\sum_p (W_{j,p} * o_p) \leq B_j$ $\quad\quad \forall j$

$E_j \leq \sum_p (W_{j,p} * l_p)$ $\quad\quad \forall j$

$D_{i,c} \leq N_j + M(2 - X_{i,j} - Y_{j,c})$ $\quad\quad \forall i, j, c \quad i \neq j$

$N_j = E_j - m_j$ $\quad\quad \forall j$

$D_{j,c} = (E_j + v_{j,c}) * y_{j,c}$ $\quad\quad \forall j, c$

$W_{j,p}, \quad X_{i,j}, \quad Z_{i,j} \quad \in \quad \{0,1\}$

$B_j, \quad T_j, \quad N_j, \quad E_j, \quad D_{j,c}, \quad \geq \quad 0$

APPENDIX E: INFORMATION TO RE-CONSTRUCT EXPERIMENTS

This appendix consists of the information we need to re-construct the same experiments for the Brynild case study and the general capabilities study.

BRYNILD CASE STUDY

We start with the parameter settings of the shift schedules. Second, we present the information of every SKU we need for the calculations to obtain the parameter values of the different products. Lastly, the production orders of every week.

Shift schedules

o_p First possible non-negative start time of processing day p (day 1 has as start time 0)
l_p Latest possible non-negative end time in hour of processing day p

Table E.0.1: Values 2-shift + Saturday schedule

***2-shift* + Saturday**	$\boldsymbol{o_p}$		$\boldsymbol{l_p}$
Monday 6:30h	0	Monday 22:30h	16
Tuesday 6:30h	24	Tuesday 22:30h	40
Wednesday 6:30h	48	Wednesday 22:30h	64
Thursday 6:30h	72	Thursday 22:30h	88
Friday 6:30h	96	Friday 14:30h	104
Saturday 6:30h	120	*Saturday 14:30h*	128

3-shift schedule starts on Sunday 22:30h indicated by 0 and end Friday 14:30h indicated by 112.

Information Intermediates

						Dynaflow or Catelli			Skap								
	Art. Nr	Produkt	Temp	temp fase 1	temp fase 2	Tid Varme	Tid Kjøl		Tid Varme	Tid Kjøl		Dyseplate	Støpe-vekt	Ant.fig	inkl. kandering	Ferdig tørk 1 rad	Ferdig tørk 2 rader
Familie A	112270/11227	Godte Gom	55			37	10	47	60	24	84	2790	0.98	800	0.78	1651	3302
	113072	Stupedama	55			38	10	48	48	24	72	2790	1.4	800	0.68	2056	4113
	113500	Påske Fruk	55			38	10	48	48	24	72	3245	6.3	168	0.84	2400	4801
	112958	Jordbærfisk	55			38	10	48	48	24	72	3245	6.4	168	0.84	2439	4877
	113482	Jul Fruktgel	55			38	10	48	48	24	72	3245	6.3	168	0.84	2400	4801
	113729	Supermixge	55			38	10	48	48	24	72	3245	6.4	168	0.84	2439	4877
	113515	Lakrisbåter	60			38	10	48	48	24	72	2792	4.1	384	0.84	3571	7141
Familie B	112815	Knatter	40-50*	40	50	48	10	58	48	24	72	2790	0.71	800	0.75	1150	2300
Familie C	107270	Frutti Bean	60	(NB! Knatter er 24t + 24 t		48	10	58	70	24	94	2791	2.5	476	0.78	2506	5012
Familie D	111916	Dent Eukaly	60			36	10	46	x	x	0	5020	1.5	660	0.68	1818	3635
	112189	Dent Trio	60	Dynaflow	Catelli	36	10	46	x	x	0	5020	1.5	660	0.68	1818	3635
	111918	Dent Salt la	60-65	65	60	48	10	58	x	x	0	5020	1.5	660	0.68	1818	3635
	113304	Dent Crush	60-65	65	60	48	10	58	x	x	0	5020	1.5	660	0.70	1871	3742
	113319	Dent Crush	60			36	10	46	x	x	0	5020	1.5	660	0.70	1871	3742
	113943	Dent Crush	60			36	10	46	x	x	0	5020	1.5	660	0.70	1871	3742
	113927	Dent Flip La	45			40	10	50	x	x	0	3245/5327	1,35 / 3,5	168	0.68	1080	2159
	113932	Dent Flip C	45			40	10	50	x	x	0	3245/5327	1,35 / 3,5	168	0.68	1080	2159
	113940	Dent Flip Jo	45			40	10	50	x	x	0	3245/5327	1,35 / 3,5	168	0.68	1080	2159
	113604	HF P.P Lak	60			36	10	46	x	x	0	2791	1.5	476	0.68	1311	2622
Familie K	105994	Gelepynt Fr	20			0	24	24	0	24	24	2791	2.4	476	0.92	2838	5675
(Kaldstøp)	107935	Jellymen ek	20			0	24	24	0	24	24	3245	9.0	168	0.93	3797	7593
	106571	Myke seigm	20			0	24	24	0	24	24	3245	6.5	168	0.93	2742	5484
	104452	Røde Hjerte	20			0	24	24	0	24	24	3245	7.8	168	0.93	3290	6581
	106497	Sure Skumf	20			0	24	24	0	24	24	3245	4.7	168	0.84	1791	3582
	113492	Figurskum-	20			0	48	48	0	48	48	3245/5327	3,7 / 6,9	168	0.78	2441	4883
	113543	Skumgele A	20			0	48	48	0	48	48	3245/5327	3,9 / 7,4	168	0.78	2618	5236
	113542	Skumgele C	20			0	48	48	0	48	48	3245/5327	3,9 / 7,4	168	0.78	2618	5236
	112873	Sure Colafl	20			0	72	72	0	72	72	3245/5327	3,7 / 6,9	168	0.87	2723	5446
	113541	Gompegelé	20			0	48	48	0	48	48	3245/5327	2,2 / 3,5	168	0.86	1365	2731
	113087	HF Skumge	20			0	48	48	0	48	48	3245/5327	3,9 / 7,4	168	0.78	2618	5236
	113717	Barnetimen	20			0	48	48	0	48	48	3245/5327	3/5,4	168	0.78	1946	3892
	113095	HF Skumge	20			0	48	48	0	48	48	3245/5327	3,9 / 7,4	168	0.78	2618	5236
	113494/113538	Søte Rakke	20			0	48	48	0	48	48	3245	5.35	168	0.85	2063	4125
	113561/113562	Sure Rakke	20			0	48	48	0	48	48	3245	5.35	168	0.93	2257	4514
	113563	Myke Rakk	20			0	48	48	0	48	48	3245/5327	3,2/5,4	168	0.85	2082	4164
Familie	111132	Sure Button	55			48	10	58	48	12	60	5136	9.6	96	0.90	2239	4479
	111892	Skumegg	48			24	10	34	24	10	34	2792	1.3	384	0.84	1132	2264
	113496	Sure tær	46			48	10	58	48	10	58	5134	6.7	156	0.90	2540	5080

Figure E.0.1: Intermediate's drying time

Production demand per week

Week 2

Intermediate	Quantity in kg
111132	6500
113543	5000
113500	7500
112189	3200
104452	8000
112958	7500
106571	8300
112815	10800
113319	4500
113542	5000

Week 3

Intermediate	Quantity in kg
111916	4500
106571	33200
112958	14000
112815	10800
113304	4500
113496	3200
113500	2450
111132	2000

Week 4

Intermediate	Quantity in kg
112873	3700
113561/113562	700
113542	700
106497	4500
113729	6500
104452	8100
113319	4500
107935	6500
111918	4500
107270	6500
112815	7200
113563	3100
113543	700
113494/113538	700

Week 5

Intermediate	Quantity in kg
113515	6200
111132	13000
106497	4500
112958	7200
112815	7200
106571	11000
111916	4500
113541	1300
113717	1500

Week 6

Intermediate	Quantity in kg
112270/112271	9200
106571	13800
112815	15300
113604	2600
112958	7200

Week 7

Intermediate	Quantity in kg
112189	3000
112958	12200
106571	7000
112815	14400
113304	3500

Week 8

Intermediate	Quantity in kg
106571	16600
111132	12600
112815	14400
112873	6100
111918	5000

TEST INSTANCES GENERAL CAPABILITIES

To test the capabilities of the MILP develop 3 types of job instances:

- Small jobs: short drying time on the parallel machines;
- Large jobs: long drying time on the parallel machines;
- Mixture of small and large jobs.

We base our instances on Brynild's input data. The parameters we determine are set-up time, cooking time, moulding time, drying time for every cabinet, and whether drying is possible in a certain drying cabinet:

- We randomly allocate set-up time that varies between 0, 1.5, and 2 hours;
- Cooking time remaines the same for every job, namely 0.5 hours;
- The average minimum moulding time of Brynild is 2.7977 hours and the average maximum moulding time of Brynild is 7.0145 hours. Therefore, we randomly select a moulding time between 2.8 hours and 7 hours for each job;
- In Brynild's case the possibility of a drying cabinet being able to dry a certain product is 85%. We use this percentage to randomly decide, which job can or cannot dry in certain drying cabinets;
- We present the drying times for the various instances in Table 0.2.

Table 0.2: Drying time for the different job instances

	Drying
Small Jobs	New cabinets= Random between 12 and 24 hours Old cabinets= 70% probability + 24 hours, 30% probability + 0 hours
Large Jobs	New cabinets= Random between 24 and 36 hours Old cabinets= 70% probability + 24 hours, 30% probability + 0 hours
Mixture of Jobs	New cabinets= 50% probability random between 12 and 24 hours, 50% probability random between 24 and 36 hours Old cabinets= 70% probability + 24 hours, 30% probability +0 hours

APPENDIX F: RESULTS

In this appendix we state some additional information of the results we find. We start with the results from the Brynild case. Thereafter, the results of the general experiments.

BRYNILD CASE STUDY

Main MILP

The main model experiments.

End-drying objective

Additional results of the end-drying objective. Both the end-drying time and the end production time in comparison to the Brynild schedules in Table F.0.1.

Table F.0.1: Results main model end-drying objective

Brynild vs MILP	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
End Drying	-20.12%	-14.92%	-26.22%	-19.46%	-21.09%	-22.5%	-19.28%
End Production	+2.08%	-1.47%	+2.02%	-2.92%	+4.76%	+2.48%	+5.87%

End Stage 1 objective

Additional results of the Stage 1 objective. Both the end Stage 1 time and the end-drying time in comparison to the Brynild schedules in Table F.0.2.

Table F.0.2: Results main model end Stage 1 objective

Brynild vs MILP	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
End Production	-19.56%	-5.13%	-25.24%	-26.45%	-0.22%	-31.52%	-0.42%
End Drying	-11.57%	-10.54%	-8.08%	-13.14%	+6.21%	-19.84%	-6.10%

Heuristics

The heuristics in comparison to the main MILP.

End-drying objective

Additional results of the end-drying objective. First the comparison of the end-drying time between the heuristics and the main model in Table F.0.3. Second, the comparison of the end production time between the heuristics and the main model in Table F.0.4.

Table F.0.3: Heuristic results end-drying objective; end-drying time

End Drying	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
Without SDST	+1.38%	+3.71%	+1.59%	+1.56%	+1.58%	+2.02%	+1.19%
Ass. & Seq. separate	+1.44%	+3.76%	+4.21%	+1.61%	+10.57%	+0.00%	+0.33%
10 min run	0.00%	+0.87%	0.00%	0.00%	0.00%	0.00%	0.00%

Table F.0.4: Heuristic results end-drying objective; end Stage 1 time

End Production	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
Without SDST	-16.03%	+5.28%	+1.93%	+1.92%	+1.87%	-19.28%	+2.19%
Ass. & Seq. separate	0.00%	+5.35%	+4.75%	+1.98%	+1.23%	-21.80%	+0.68%
10 min run	-17.50%	-0.60%	0.00%	0.00%	0.00%	0.00%	0.00%

End Stage 1 objective
Additional results of the Stage 1 objective. First the comparison of the end production time between the heuristics and the main model in Table F.0.5. Second, the comparison of the end-drying time between the heuristics and the main model in Table F.0.6.

Table F.0.5: Heuristic results end Stage 1 objective; end Stage 1 time

End Production	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
Without SDST	+6.00%	+1.62%	+3.66%	+5.71%	+0.53%	+17.62%	+0.83%
Ass. & Seq. separate	+6.48%	+0.81%	+3.59%	+0.83%	+0.53%	0.00%	+0.42%
10 min run	+1.47%	+3.04%	+1.22%	+3.36%	0.00%	0.00%	0.00%

Table F.0.6: Heuristic results end Stage 1 objective; end-drying time

End Drying	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8
Without SDST	+3.30%	+5.65%	-2.75%	+2.30%	-15.85%	+1.21%	+7.04%
Ass. & Seq. separate	+0.18%	+4.37%	-8.35%	+0.96%	-16.33%	0.00%	-0.64%
10 min run	+0.81%	+10.56%	0.00%	+11.08%	0.00%	0.00%	0.00%

General experiments
The additional information of the general experiments regarding the different models. We present the comparison between the objective values for the optimal schedule and the objective values when we stop the scheduling after 5 minutes, see Figure F.0.1, Figure F.0.2, and Figure F.0.3.

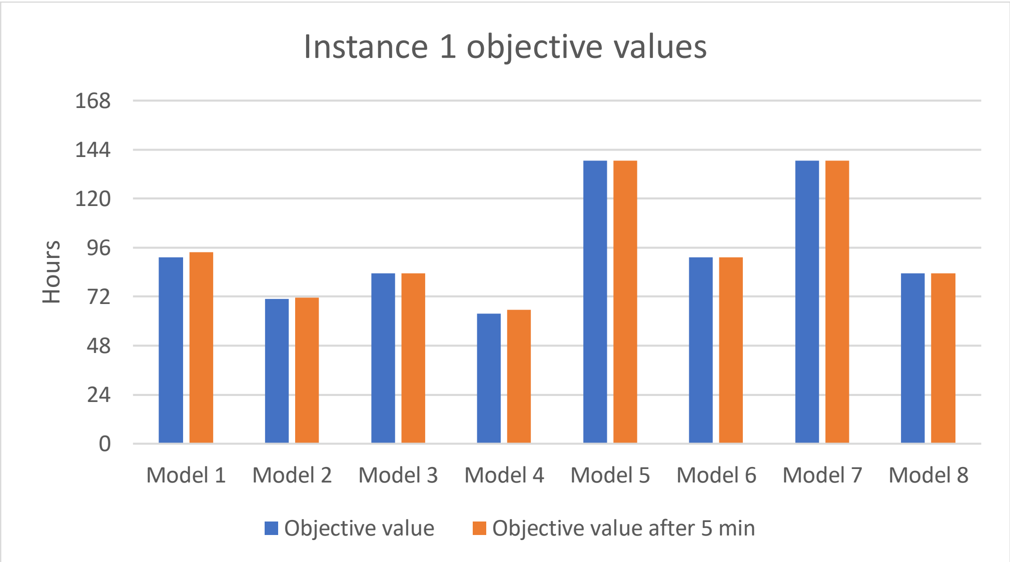

Figure F.0.1: Objective values instance 1

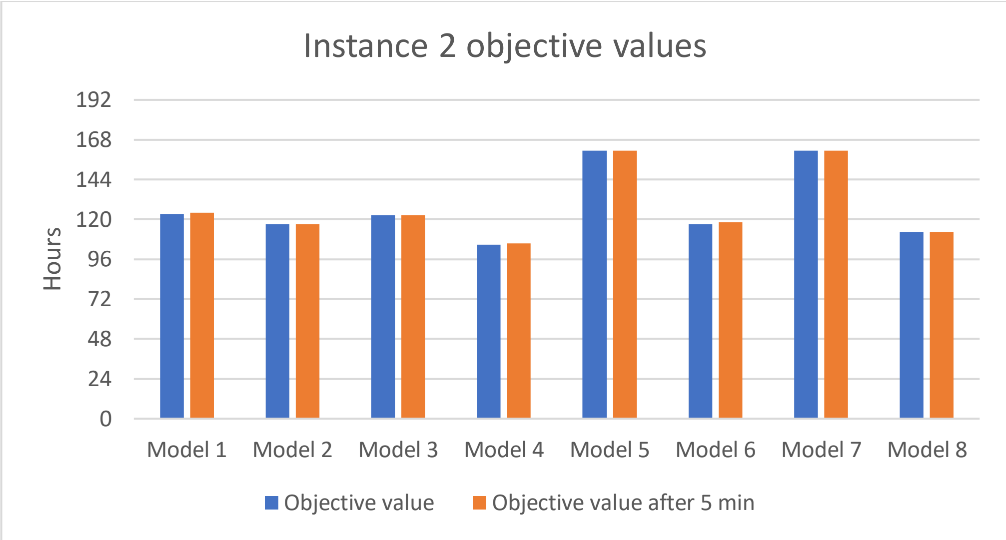

Figure F.0.2: Objective values instance 2

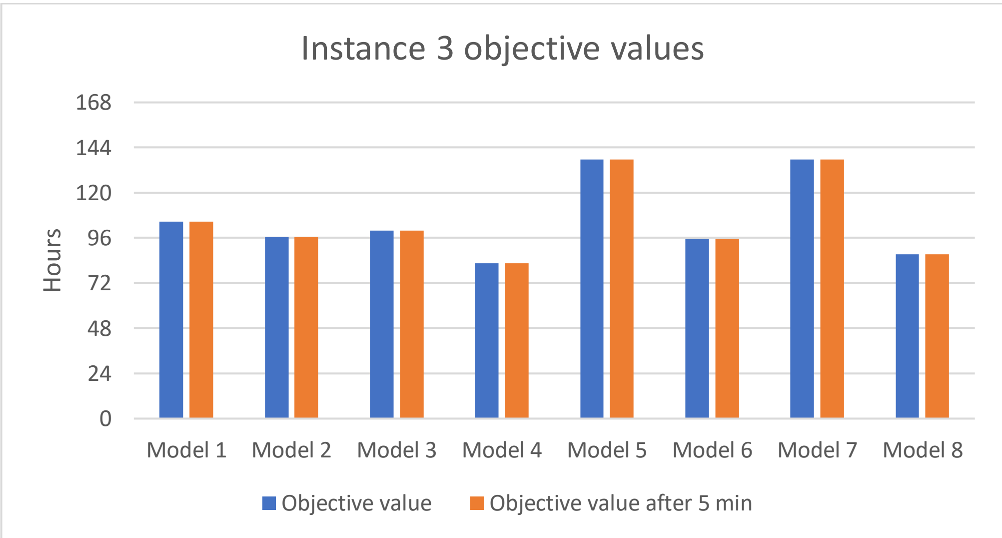

Figure F.0.3: Objective values instance 3

APPENDIX G: EVALUATION RESULTS

This appendix consists of additional evaluation regarding results from Chapter 5.

PAIRED T-TESTS OBJECTIVE VALUES

End-drying time objective

To be certain that the reduction in time is significant we perform a paired t-test (Kim, 2015). We use the following hypothesis:
H_0: There is no significant change, on average, in the end drying outcome of Brynild and our MILP.
H_1: There is an average significant change in drying end time.

We use α= 0.05 for the level of significance and find a p-value of 0.000005339 for the one-tail. We use the one-tail test as we know beforehand that all our MILP values are lower than the Brynild's values. However, even with a two-tail the p-value is 0.00001068. Therefore, we reject H_0 hypothesis with a significance level of 0.05 and we state that the differences in end-drying time is significant.

Stage 1 end time objective

To be certain that the reduction in time is significant we perform a paired t-test (Kim, 2015). We use the following hypothesis:
H_0: There is no significant change, on average, in the end Stage 1 outcome of Brynild and our MILP
H_1: There is an average significant change in Stage 1 end time

We use α= 0.05 for the level of significance and find a p-value of 0.01195 for the one-tail. Also for the Stage 1 objective we use the one-tail test as we know beforehand that all our MILP values are lower than the Brynild's values. However, even with a two-tail the p-value is 0.0239. Therefore, we reject H_0 with a significance level of 0.05 and we state that the differences in end-drying time is significant.

EXAMINATION OF SCALING DOWN RESULTS

From Table 5.2, three out of the four 3-shift schedules we can schedule using a 2-shift + Sat except for Week 3. One of the three 2-shifts + Saturday schedules we can schedule using a 2-shift schedule. The difference between a 3-shift schedule, 2-shift schedule + Saturday and a 2-shift schedule is the available production time in Stage 1. The amount of drying time remains the same. To research the reason why we can reduce certain weeks and others not, we examine the constraint that determines the time needed for Stage 1.

$$T_j = k_j + m_j + \sum_{i \in J}(s_{i,j} * Z_{i,j}) \qquad \forall j \qquad [4.1.8]$$

The production time is a variable due to the sequence dependent set-up times. Therefore, we analyse different production times. The first option is without any set-up times, so the minimum the production time certainly needs. The second option is the production time with the average of the set-up times of that week and the third option is the production time, which uses 2 hours as set-up, the maximum production time. The calculations for each week we present in Table G.0.1.

Table G.0.1: Stage 1 production hours per week

Week	Original Brynild schedule	Min production hours	Average production hours	Max production hours
2	3-shift	60.73 h	79.00 h	80.73 h
3	3-shift	67.82 h	87.59 h	91.82 h
4	3-shift	59.70 h	78.33 h	79.70 h
5	3-shift	53.29 h	71.75 h	73.29 h
6	2-shift + Sat	53.07 h	66.87 h	71.07 h
7	2-shift + Sat	43.83 h	56.38 h	59.83 h
8	2-shift+ Sat	53.15 h	67.95 h	71.15 h

Before examining the information we provide in Table G.0.1, we need the production hours maximum available in the different shift schedules, see Table G.0.2.

Table G.0.2: Production hours available

Shift schedule	Stage 1 hours available
3-shift	112 hours
2-shift + Saturday	80 hours
2-shift	72 hours

If the problem is a two-stage HFS with infinite buffers between the 2 stages, we only have to determine whether the maximum production hours is less than the production hours available in Stage 1. What means that we can schedule all 2-shift + Saturdays schedules within in a 2-shift schedule.

Even with the information above, it is still difficult to predict whether a week worth of production fits in a certain shift week without running the model. We do not give this inside with much certainty as the test population is small. However, we do suspect that to scale down the 3-shift schedules to a 2-shift + Saturday it is necessary for the average production hours to be under 80 hours. This assumption is not farfetched as the additional Saturday gives room for drying cabinets to finish, and fill again within that time. The available production time on Stage 1 becomes less, however better spread over the week, which the schedule needs to finish drying and make place for the next job. We also see that in in the 3-shift schedule there is a lot of vacant time. Even for Week 3, Stage 1 is not in use for 112-91.92 = 20 hours.

The statements above are all under the assumption that the drying time of the jobs is based on Brynild's indicated drying times. When the drying times become way longer, it is obviously not possible to produce everything in 1 week. So, these statements are only valid when the 3-shift schedule is already proven feasible.

To be able to go from a 2-shift + Saturday to a 2-shift schedule, it is not enough to have the average production time under the production time available as we observe for Week 6 and Week 8. When we remove the Saturday from the schedule, an important additional day for the drying cabinets to finish drying and be refilled is absent. We do not have enough information to do an estimated guess of how much time we should be under the 72 hours. We suspect that going from a 2-shift + Saturday to a 2-shift schedule is influenced more by the available drying time and the drying time needed. We do not examine this any further.