

A decision support system for the operational production planning and scheduling of an integrated pulp and paper mill

Gonalo Figueira^{a,*}, Pedro Amorim^a, Lu s Guimar es^a, M rio Amorim-Lopes^a, F bio Neves-Moreira^a, Bernardo Almada-Lobo^a

^aINESC TEC, Faculdade de Engenharia, Universidade do Porto, Porto, Portugal

Abstract

Production planning and scheduling in the process industry in general and in the pulp and paper (P&P) sector in particular can be very challenging. Most practitioners, however, address those activities relying only on spreadsheets, which is time-consuming and sub-optimal. The literature has reported some decision support systems (DSSs) that are far from the state-of-the-art with regard to optimization models and methods, and several research works that do not address industrial issues. We contribute to reduce that gap by developing and describing a DSS that resulted from several iterations with a P&P company and from a thorough review of the literature on process systems engineering. The DSS incorporates relevant industrial features (which motivated the development of a specific model), exhibits important technical details (such as the connection to existing systems and user-friendly interfaces) and shows how optimization can be integrated in real world applications, enhanced by key pre- and post-optimization procedures.

Keywords: Decision Support System, Lot-sizing and Scheduling, MIP-based heuristics, Pulp and Paper Industry, Continuous Production

1. Introduction

The pulp and paper (P&P) industry is highly capital intensive, which means that investments in capacity can represent very long-term decisions. For instance, the modification of a single paper machine requires a planning horizon of at least five years (Martel et al., 2005). Paper products are commodities, with their price being determined in the market and characterized by small margins. Hence, companies must differentiate themselves by improving customer satisfaction indicators, while keeping production costs as low as possible. Furthermore, the P&P production process is also energy intensive. Producing one tonne of paper requires 5–17 GJ of process heat, depending on the type of paper and on the technology applied (Szab  et al., 2009). The P&P industry uses 84% of the fuel energy consumed by the forest products industry as a whole and is one of the largest producers of greenhouse gas emissions (Jankunaite, 2006). Therefore, it is of particular interest in the context of environmental discussions.

These three main factors (capital intensity, energy intensity and competitive market) make the production planning an essential activity in the quest for improvements in operational efficiency

*Corresponding author. Tel.: +351 222094364. E-mail address: goncalo.figueira@fe.up.pt (G.Figueira).

and consequently economic gains. However, the planning process poses a variety of challenges, both to practical and scientific fields (e.g. combinatorial nature of paper grades scheduling, shifting bottlenecks, variable production rates, etc. – see Section 2). If these challenges are successfully addressed, companies can achieve a true competitive advantage.

In most cases, production planning is addressed manually by practitioners, even in modern mills with sophisticated automated systems. That applies not only to the P&P industry, but to the process industry in general (see Harjunkoski et al. (2014)). Companies may use advanced tools for particular tasks, such as the planning of the paper reels' cutting, but when addressing the overall planning activity (e.g. size and sequence of paper campaigns, production rates, etc.) most practitioners rely only on spreadsheets. This manual process is time-consuming, sub-optimal (as only few alternatives are considered) and completely dependent on the planners' expertise.

Therefore, there is the need for optimization-based tools that support decision-making in the operational production planning of these mills. Some decision support systems (DSSs) for this particular industry were reported in the literature (e.g. Murthy et al. (1999); Respício and Captivo (2008)). Although these systems consider relevant industrial features and practical issues, the underlying approaches are far from the state-of-the-art in production planning, as they are based on simplistic models and heuristics. On the other hand, more research-oriented work does not address the issues of an industrial implementation.

Even in other process industries, the literature on the implementation of this type of DSSs is scarce. Bongers and Bakker (2006) describe a scheduling problem in a medium size ice cream plant. However, the model has not been validated or compared against manual schedules. This problem was subject to further studies and improvements/extensions (Kopanos et al., 2012; van Elzакker et al., 2012), but it seems that they have not been implemented in practice. Janak et al. (2006) and Shaik et al. (2009) address production scheduling problems in large-scale batch and continuous plants, respectively, of a chemical company. Nevertheless, these papers have not reported any industrial implementation. Wassick (2009) is one of the few comparing manual schedules to schedules generated by optimization methods. The author has approached a variety of optimization problems in an integrated chemical complex. The problems included production planning, energy system scheduling, site reliability design and waste treatment network scheduling. For the same chemical company, Wassick and Ferrio (2011) have developed a generic formulation which they implement in a custom interface within Microsoft Excel and apply to different optimization problems, including storage management, production scheduling and container loading.

Our work helps to close the gap between research and practice, proposing an optimization-based DSS, which improves on manual planning and contributes to the literature in the following ways:

- exploring desirable characteristics of analytical models and methods;
- identifying relevant industrial features and how they should be included in the solution approach;
- extending models of the literature, considering practical constraints and objectives;
- exhibiting important data processing in both pre- and post-optimization phases;
- illustrating required connections to existing systems and desirable aspects in the user interface.

We start by describing the industrial system and the planning challenge in Section 2. This motivates the discussion in Section 3 on the different approaches to operational production planning in the literature. In that section we explore some desirable characteristics of optimization models and identify relevant features for industrial practice. Based on that, we choose an appropriate mathematical model and extend it in order to include the practical issues identified. Our formulation is presented in Section 4. Section 5 motivates and explores the solution method for an efficient yet simple and flexible resolution of this complex problem. Section 6 details the integration of the optimization in the decision support system, describing pre- and post-optimization steps, as well as the interfaces to the existing information systems and the final user. The usage of the DSS is also discussed. In Section 7 we compare a plan obtained with the DSS to one generated manually by production planners, and provide further details on the performance and usage of the DSS at the company. Finally, the last section summarizes the benefits of our DSS, discusses its applicability to other environments and shows possible directions for further improvement. Since the paper touches different fields and a variety of concepts, a complete list of acronyms is provided in Appendix A. The full mathematical notation is condensed in Appendix B.

2. The challenge

The P&P production process is illustrated in Figure 1. The variables depicted in this figure will be introduced later in Section 4. In a first step, both virgin and recycled pulps are produced out of wood and recycled paper, respectively. These pulps are then stored in tanks, waiting to be pulled by the paper machine. The machine can produce different types of paper. Each type of paper (or grade) is characterized by its grammage (measured in g/m^2) and pulp mixture. The configuration of the machine to produce a different grade is sequence-dependent, for instance, changing from $170 g/m^2$ to $200 g/m^2$ is considerably less costly than from $115 g/m^2$ to $200 g/m^2$. Each setup leads to a loss in the production process in terms of time and quantity of a lower quality paper produced (as the machine is never idle and the paper produced will not be homogeneous, nor completely satisfy the customer requirements). The wasted paper (setup loss) is dissolved and stored in the loss pulp tank, to be pulled again by the paper machine.

The master reel that results at the end of the paper machine, the jumbo, is cut into smaller reels. The paper wasted in the cutting stage (trim loss) is also fed back to the production process. Customers place orders for reels of different widths and grades. The orders may have different priority levels. The maximum priority is given to those that travel by ship, since the company has to schedule containers in advance and commit to a given due date. Then, the remaining is divided into normal and priority orders.

In parallel, a by-product of the digester's pulping process, the weak black liquor, is concentrated in evaporators and burnt in the recovery boiler to provide high-pressure steam and to regenerate the chemicals applied in the pulping stage. The steam can either be used for the paper drying process or be led to counter-pressure turbines which produce electrical energy to be sold afterwards.

In other companies, the paper mill can be physically distant from the pulp and recovery plants. However, integrated plants, like that of the case study, represent 65% of the industry (CEPI, 2013) and are more capable of achieving high levels of both energy and economic efficiency, due to:

- energy conservation (e.g. direct use of steam in the paper drying process);




Figure 1: The integrated pulp, paper and recovery plant, with multiple distribution channels.

- absence of additional processes (e.g. pulp drying);
- material closed loops (e.g. recycling of paper machine setup and trim losses) which make it possible to reduce waste production and energy;
- tightly integrated equipment, which reduces the required capacity.

Nevertheless, an integrated mill poses additional difficulties. The complex production process described above (with convergent, divergent and loop flows) and the tightly integrated equipment, with limited intermediate storage space, result in multiple and shifting bottlenecks. For instance, the company under study produces two main products (KLB and VLB), which have major differences with respect to the incorporation of virgin and recycled fibres. Hence, the bottleneck stage clearly shifts according to the mixture being produced (for example, VLB quickly exhausts the recycled pulp mill, whereas KLB is typically restricted by the recovery boiler).

Moreover, the orders placed by customers put a great pressure on production, as the system operates in a make-to-order (MTO) policy. Still, adjusting the production sequence to better meet market needs has to be done carefully, since the desired production cycles will be disrupted, and some stages may be forced to drastically reduce their production rate or even to shut-down. The start-up of the process after these interruptions is typically problematic and requires a large expenditure of energy, increasing the environmental load of the mill. The rates of the various resources need to be as steady as possible, as rough changes can cause quality deviations and undesired wear of equipment. Therefore, the variation that results from production cycles is conveyed as much as possible to the tank levels. Nonetheless, the system is subjected to process variability (for instance, in production yield, required pulp mixture, etc.) and disturbances (such as paper breaks, equipment failures, etc.). Thus, it is important to keep slacks in the intermediate buffers, in order to prevent under and overflows.

In this capital intensive industry in general, and in this plant in particular, exploiting existing capacity is a major concern. Therefore, the plant works on a 24/7 basis and backlog is tolerable. When accepting orders, managers use a base sequence of grades, where they can verify the earliest available slot for that grade and hence commit to a due date (available-to-promise). The base sequence minimizes changeovers, since not only they consume time (i.e. capacity), but can also be highly disruptive to the system, breaking the production stability and causing wear to




Figure 2: Example of a production plan – cycle of grades (left axis) and cycle of virgin pulp tank, measured in tonnes (right axis).

equipment. Nevertheless, at the operational level, the sequence must be adapted according to the evolving state of the system.

The operational planning consists of determining the sequence and size of paper campaigns (where each campaign represents one of the 29 commercial grades), and all the production rates of the different resources (paper machine, pulp digester, recycled pulp mill, evaporators and recovery boiler) for the following 10 to 15 days. This horizon is important to effectively check production cycles and the impact on every unit. An example of a production plan is illustrated in Figure 2. The cycle of grades is conveyed to a cycle of the virgin pulp tank level (as the digester's rate should be as smooth as possible). As VLB-type grades consume more recycled pulp, the tanks of virgin pulp increase their stock throughout the VLB campaigns. This stock is consumed during most KLB-type campaigns. The slacks of the tanks are kept to some desired limits that will not impact the feasibility of the subsequent plans.

This planning process is conducted by two main entities: Sales and Production departments. The former aims to fulfil the orders respecting due dates, whereas the latter is focused on operational costs and throughput maximization. As stated above, aligning these two main objectives is problematic. Therefore, when devising a plan, several interactions are needed between both departments. The lack of a holistic view and of an optimization-based DSS makes the generation of feasible plans a difficult task and their quality is compromised. This is especially true in the presence of disturbances, scheduled maintenance or rush customer orders. The latter prevent the reels cutting from being planned for the same horizon. Indeed, one late order can completely spoil the cutting patterns. Therefore, managers choose to define those patterns just for the following 2 to 3 days.

3. Solution approach

Considering the aforementioned challenges to the production planning activity, we propose an optimization-based DSS. Our system supports decision-making regarding the size and sequence of paper campaigns, the production rates for every unit and the fulfilment of individual orders (with different priority levels). The planning of cutting patterns is not included for three

reasons: (i) companies already have software packages for that task; (ii) the planning horizon is radically different, as the cutting is only scheduled for a couple of days ahead; (iii) because there are buffer areas between the paper machine and the cutting reels, jumbo storage is allowed to some extent, making it possible to decouple the two processes.

However, at this planning level it is important to integrate lot-sizing and scheduling. Indeed, the sequence dependent setups in the paper machine, with important costs related to production stability, must be taken into account when defining production plans. This is not possible in plain lot-sizing models. On the other hand, pure scheduling approaches are not appropriate either, since this (capital intensive) continuous production plant has high utilization ratios. Therefore, choosing the amounts to be produced and the respective sequence (leaving orders partially or totally unmet or to be met later) is a key decision.

The literature on this industry has mainly focused on either lot-sizing (e.g. Rizk et al. (2008), Poltroniere et al. (2008)) or scheduling (e.g. Akkiraju et al. (2001), Correia et al. (2012)). Moreover, these studies have not included the pulping and recovery stages. This is crucial in integrated plants due to the tightly integrated equipment, and the existence of multiple and shifting bottlenecks. Santos and Almada-Lobo (2012) were the first to integrate those stages. The proposed model was based on the general lot-sizing and scheduling problem (GLSP), presented by Fleischmann and Meyr (1997), and considered sequence dependent setup times and costs. Figueira et al. (2013) then improved the resolution of the problem with a variable neighbourhood search, whereas Furlan et al. (2015) developed a genetic algorithm for the multi-machine scenario.

GLSP is a small-bucket model (also known as time slot formulation), i.e., allows just one setup in each (micro) time period. Having only one product in each period is a key aspect in our case, since it allows the accurate evaluation of mass balances. This is crucial when the intermediate tanks have tight capacities. In big-bucket models (also known as precedence-based), multiple products with different rates are produced in the same period, whereas the mass balances are only checked at the end of each period. Therefore, the tank limits may be violated within a period, while being respected at the beginning and at the end. Verifying tank levels during each period would not be impossible in big-bucket models, but it would certainly be more difficult, and could require some simplifying assumptions. Still, these models are recommended if such constraints are not critical, since they are considerably more efficient (Amorim et al., 2013). These models, usually in the form of the so-called CLSD (capacitated lot-sizing with sequence dependent setups), have been applied to similar process industries (Almada-Lobo et al., 2010; Kopanos et al., 2011b,a).

Contrary to other small-bucket models (such as the discrete lot-sizing and scheduling problem DLSP (Fleischmann, 1994)), GLSP combines continuous- and discrete-time grids. Continuous-time grids are important to accurately track the production of non-standard amounts, essential in this MTO system. Discrete-time grids make it possible to assess the fulfilment of demand (at discrete time points) during the planning horizon. This mixed grid thus makes the interface between Production and Sales requirements.

Other works in the literature, in both the operations research (OR) and process systems engineering (PSE) communities, sought to combine the advantages of continuous and discrete-time grids. Maravelias (2005) and Westerlund et al. (2007) proposed models that use a discrete grid, but where campaigns do not need to start or end at the exact grid points. This aspect was also explored by the continuous setup lot-sizing problem (CSLP), presented by Karmarkar and Schrage (1985), and the proportional lot-sizing and scheduling problem (PLSP), proposed by Drexel and Haase (1995). However, obtaining continuous amounts in these models implies idle times and/or multiple products per period, which spoils the evaluation of intermediate tank levels.




Figure 3: GLSP – each period t is divided into multiple continuous slots, where a grade j is produced.




Figure 4: CGLSP – the discrete grid (of periods t) is crossed by continuous slots, where a grade j is produced.

Models based on GLSP have been extensively explored in the literature (e.g. Lim and Karimi (2003), Erdirik-Dogan and Grossmann (2008), Castro et al. (2009), Amorim et al. (2011)). Guimarães et al. (2014b) compare different GLSP formulations and show that the model proposed by Clark and Clark (2000) performs best in the presence of minimum lot constraints. Still, all these models divide each (macro) period into multiple slots (micro-periods). Thus, if a campaign crosses a period, it will require more than one slot (see Figure 3). This increases the difficulty of incorporating some practical constraints. For instance, in our case we need to fix the amounts of the first campaigns, as they were already programmed for the cutting stage. In those models we have to determine the number of slots for each of these campaigns. Since the production rates are variable, and will have to be adapted according to the state of the system, the appropriate number of slots can only be approximately estimated. This may compromise the feasibility of the plans generated. Other constraints include maximum and minimum lot sizes and a minimum distance between campaigns of the same grade.

Karimi and McDonald (1997) and Camargo et al. (2012) proposed other mixed grids formulations, where (continuous) production slots may cross (discrete) time periods (see Figure 4). This feature allows assigning only one slot to each campaign, and hence easily handle the aforementioned issues. The authors concluded, however, that it is more difficult to solve these crossed grids models (here called CGLSP) than the GLSP. There is thus a trade-off between modelling accuracy and performance. For our P&P industrial problem, the CGLSP was found to be the most appropriate choice. In fact, the flexibility in incorporating the numerous practical constraints (cf. Subsection 4.6), avoiding specific heuristics for the allocation of slots, was crucial in this industrial implementation, as feasibility was occasionally compromised with GLSP.

In addition to the multi-stage environment and the sequence dependent setup times and costs modelled by Santos and Almada-Lobo (2012), our system comprises a set of features that are important for industrial practice, such as:

- fixing the size of certain campaigns (due to cutting programs);
- fixing the sequence of certain campaigns, critical when the production system is not responding well to certain grades;
- giving an initial sequence of campaigns, which can then be improved;
- scheduling stoppages in the paper machine and digester (for maintenance purposes or in a disturbance situation);
- variable production rates for all the units, including the paper machine;

- smoothness of production rates (to be performed in post-optimization, in order to keep model's linearity);
- different priority levels for client orders;
- demand beyond the planning horizon (important in low demand periods).

The DSS was designed to support both Sales and Production staff, weighting their conflicting objectives. Other studies in the literature (e.g. Murthy et al. (1999)) proposed multi-objective approaches, where managers can evaluate different plans and choose the one that seems more appropriate. However, in our case this evaluation appears to be time consuming. Indeed, at the operational level managers do not have the time to properly evaluate different plans. Therefore, they prefer to define *a priori* the trade-offs between objectives, according to the company's policies, and then use them in every plan generation.

4. Mathematical model

We present the mathematical formulation in four main parts:

- continuous time slots and their intersection with the discrete time grid;
- production system, namely the P&P production;
- demand fulfilment, discriminated by time period;
- objective function, with all the components of the production and marketing objectives;

An illustrative example is provided to help the understanding of the model. Two additional sections are then presented to demonstrate the applicability of the model, namely:

- valid inequalities that help to improve the lower bound of the formulation;
- practice issues, such as steady production rates (essentially performed in post-optimization), undesired sequence of campaigns and stoppages.

4.1. Time slots

The production slots' lengths are variable and independent of the discrete time grid (they may cross one or multiple periods – see Figure 4). Each slot contains only one paper campaign and starts always with the changeover from the grade produced in the previous slot to the grade of the current slot. This continuous-time representation prevents time slots from being split, which will greatly facilitate the incorporation, in a straightforward manner, of industrial features, such as lower and upper bounds on the quantity produced. On the other hand, (external) demand is given for each (discrete) time period, and therefore there is a need to define the quantity produced in each of those time periods. The following set of constraints assure the coherence of this two-level time grid.

Indices and parameters

t	Index for period ($t \in [T] = \{1, \dots, T\}$)
s, u	Indices for slot ($s, u \in [S] = \{1, \dots, S\}$)
j, k	Indices for paper grades ($j, k \in [K] = \{1, \dots, K\}$)
cap	Period's capacity (in hours)
$hcap$	Planning horizon's capacity (in hours): $hcap = T \cdot cap$
$slots_{max}$	Maximum number of slots crossing one period

Decision variables

S_s^{start}	Starting time of slot s
S_s^{end}	Ending time of slot s
U_{js}^{start}	Starting time of slot s for grade j
U_{js}^{end}	Ending time of slot s for grade j
E_{ts}	$\begin{cases} 1 & \text{if slot } s \text{ intersects period } t \\ 0 & \text{otherwise} \end{cases}$
Y_{js}	$\begin{cases} 1 & \text{if the paper machine is set up for grade } j \text{ in slot } s \\ 0 & \text{otherwise} \end{cases}$
N_s	Length of slot s (in hours)

The maximum number (S) of slots is given. The planning horizon starts with the first slot ($S_1^{start} = 0$) and ends with the last $S_S^{end} = hcap$. In case the slots are not all used, empty slots can be moved to the end of the horizon as it will be shown later.

$$S_s^{start} = S_{s-1}^{end}, s \in [S] \setminus \{1\} \quad (1)$$

$$N_s = S_s^{end} - S_s^{start}, s \in [S] \quad (2)$$

$$U_{js}^{end} - hcap \cdot Y_{js} \leq U_{js}^{start} \leq U_{js}^{end}, j \in [K], s \in [S] \quad (3)$$

$$S_s^{start} - hcap \cdot (1 - Y_{js}) \leq U_{js}^{start} \leq S_s^{start} + hcap \cdot (1 - Y_{js}), j \in [K], s \in [S] \quad (4)$$

$$S_s^{end} - hcap \cdot (1 - Y_{js}) \leq U_{js}^{end} \leq S_s^{end} + hcap \cdot (1 - Y_{js}), j \in [K], s \in [S] \quad (5)$$

Constraints (1) prevent idleness by assigning the starting time of a slot to the ending time of the preceding one. The length of each slot is computed in (2). The production time of each campaign is limited to the respective bucket – see (3)–(5). When these constraints are active ($Y_{js} = 1$), clearly $U_{js}^{start} = S_s^{start}$ and $U_{js}^{end} = S_s^{end}$.

It is important to recall that the production resources have slots which are independent of the time period's boundaries. Now, the intersection between slots and time periods is assured by variables E . Naturally, the first slot takes place in time period one ($E_{11} = 1$), and the last in the last period ($E_{TS} = 1$).

$$slots_{max} \cdot T \cdot (1 - E_{ts}) \geq \sum_{t'=1}^{t-1} \sum_{s'=s+1}^S E_{t's'}, t \in [T] \setminus \{1\}, s \in [S] \setminus \{S\} \quad (6)$$

$$E_{t-1,s} + E_{t+1,s} \leq E_{ts} + 1, t \in [T] \setminus \{1, T\}, s \in [S] \quad (7)$$

The proper assignment of slots to time periods is established by (6). Because a slot s is assigned to period t , a subsequent slot cannot be assigned to any precedent time period, which guarantees

that the slots are in the correct order. If the slot spans from period $t - 1$ to period $t + 1$, then (7) ensures that period t is also crossed by the same slot.

4.2. Production system

In this subsection we focus on the pulp plant and paper mill of the P&P manufacturing system, although all the main production stages (except the reels cutting) are addressed in our DSS (see Figure 1).

Parameters

V_{min}^{dig} (V_{max}^{dig})	Minimum (maximum) rate of the digester (in rpm)
α	conversion factor of digester's throughput (from rpm to tonnes per hour)
I_{min}^{virg} (I_{max}^{virg})	Minimum (maximum) level of virgin pulp stocked in the tanks (in tonnes)
f_j	Percentage of water in paper grade j
tlb_j	Average trim loss (in percentage) of paper grade j that goes to the loss pulp tank
b_j^{virg}	Percentage of virgin pulp used in the production of paper grade j
b_{loss}^{virg}	Percentage of virgin pulp in the loss pulp
sl_{kj}	Paper lost in a changeover from grade k to j (in tonnes)
st_{kj}	Time lost in a changeover from grade k to j (in tonnes)
p_j^{min}	minimum time (at machine's maximum rate) to produce one tonne of paper grade j (in hours)
p_j^{max}	maximum time (at machine's minimum rate) to produce one tonne of paper grade j (in hours)

Decision variables

X_s^{virg}	Production of virgin pulp in slot s (in tonnes)
O_s^{virg}	Amount of virgin pulp fed into the paper machine in slot s (in tonnes)
I_s^{virg}	Inventory of virgin pulp at the end of slot s (in tonnes)
X_s^{loss}	Production of loss pulp in slot s (in tonnes)
O_s^{loss}	Amount of loss pulp used in slot s (in tonnes)
X_{jts}	Production of paper grade j in period t and slot s (in tonnes)
Z_{kjs}	$\begin{cases} 1 & \text{if a changeover from grade } k \text{ to } j \text{ occurs at the beginning of slot } s \\ 0 & \text{otherwise} \end{cases}$

In terms of pulp production, the amount of virgin pulp produced by the digester in each slot is bounded by its speed limits, according to (8). Produced pulp is stocked in tanks and then fed into the paper machine (9). In (10), this stock is also limited by lower and upper bounds. The reader is referred to Santos and Almada-Lobo (2012) for a thorough analysis of the constraints of these two areas.

$$\alpha \cdot V_{min}^{dig} \cdot N_s \leq X_s^{virg} \leq \alpha \cdot V_{max}^{dig} \cdot N_s, \quad s \in [S] \quad (8)$$

$$X_s^{virg} + I_{s-1}^{virg} = O_s^{virg} + I_s^{virg}, \quad s \in [S] \quad (9)$$

$$I_{min}^{virg} \leq I_s^{virg} \leq I_{max}^{virg}, \quad s \in [S] \quad (10)$$

The requirements of the recycled pulp and recovery plant are similar to those from the virgin pulp plant, and therefore will not be presented here (such as the minimum and maximum production of recycled pulp, liquors and steam; the inventory balance for recycled and loss pulps and liquors; and the inventory limits for recycled and loss pulps and liquors).

$$\sum_j Y_{js} = 1, s \in [S] \quad (11)$$

$$\sum_k Z_{jks} = Y_{j,s-1}, j \in [K], s \in [S] \quad (12)$$

$$\sum_k Z_{kjs} = Y_{js}, j \in [K], s \in [S] \quad (13)$$

$$X_{jts} \leq cap \cdot Y_{js}, j \in [K], t \in [T], s \in [S] \quad (14)$$

$$p_j^{min} \cdot X_{jts} \leq cap \cdot E_{ts}, j \in [K], t \in [T], s \in [S] \quad (15)$$

$$\sum_j b_j^{virg} \cdot (1 - f_j) \cdot \left(\sum_t X_{jts} + \sum_k sl_{kj} \cdot Z_{kjs} \right) = O_s^{virg} + b_{loss}^{virg} \cdot O_s^{loss}, s \in [S] \quad (16)$$

$$X_s^{loss} = \sum_j \sum_k (1 - f_j) \cdot sl_{kj} \cdot Z_{kjs} + \sum_j \sum_t (1 - f_j) \cdot tlb_j \cdot X_{jts}, s \in [S] \quad (17)$$

$$\sum_t p_j^{min} \cdot X_{jts} + \sum_k st_{kj} \cdot Z_{kjs} \leq U_{js}^{end} - U_{js}^{start}, j \in [K], s \in [S] \quad (18)$$

$$\sum_t p_j^{max} \cdot X_{jts} + \sum_k st_{kj} \cdot Z_{kjs} \geq U_{js}^{end} - U_{js}^{start}, j \in [K], s \in [S] \quad (19)$$

When producing paper, from (11) exactly one grade must be produced in each slot. Constraints (12) and (13) link the grade changeover variables to the setup state variables on the paper machine. A grade j is only produced in a given bucket in case the machine has been appropriately set up (14). Naturally, variable X_{jts} only takes positive values in case slot s crosses period t , as ensured by equations (15). Note that the production quantity (X_{jts}) is described per period, i.e., the portion of slot s that crosses period t is taken into account to define the quantity that meets the demand for that period. The virgin pulp used by the paper machine, defined in (16), comes from the virgin pulp stored in tanks, together with (part of) the pulp that is recovered from the production losses due to the grade setup changeovers and trim sub-process, which are computed in (17).

Finally, according to (18)-(19) the time confined to the setup and production operations of each grade on the paper machine in each slot is given by the length of the slot, where the machine operates within its minimum and maximum speed limits. Note that in case grade j is not assigned to slot s , from (3), $U_{js}^{end} = U_{js}^{start}$, and therefore no operations take place. The definition of the consumption of recycled pulp is modelled similarly to the consumption of the virgin pulp presented before, and for that reason it is not presented here.

4.3. Demand fulfilment

Demand is met at the end of the discrete periods. The following equations look at the intersection of the continuous production slots with those discrete periods.

Parameters

d_{jt}	Demand for paper grade j in period t (in tonnes)
tl_j	Average trim loss (in percentage) of paper grade j
I_{j0}	Initial inventory of paper grade j (in tonnes)
IB_{j0}	Initial backlog of paper grade j (in tonnes)

Decision variables

W_{jts}^{start}	Starting time of slot s for grade j in period t
W_{jts}^{end}	Ending time of slot s for grade j in period t
I_{jt}	Inventory of paper grade j at the end of period t (in tonnes)
IB_{jt}	Backlog of paper grade j at the end of period t (in tonnes)

$$P_j^{min} \cdot X_{jts} \leq W_{jts}^{end} - W_{jts}^{start}, j \in [K], t \in [T], s \in [S] \quad (20)$$

$$P_j^{max} \cdot X_{jts} \geq W_{jts}^{end} - W_{jts}^{start}, j \in [K], t \in [T], s \in [S] \quad (21)$$

$$W_{jts}^{end} \leq \min(t \cdot cap; U_{js}^{end} + hcap \cdot (1 - E_{ts})), j \in [K], t \in [T], s \in [S] \quad (22)$$

$$W_{jts}^{start} \geq \max\left((t-1) \cdot cap; U_{js}^{start} + \sum_k st_{kj} \cdot Z_{kjs} - hcap \cdot (1 - E_{ts})\right), j \in [K], t \in [T], s \in [S] \quad (23)$$

$$\sum_t (W_{jts}^{end} - W_{jts}^{start}) + \sum_k st_{kj} \cdot Z_{kjs} = U_{js}^{end} - U_{js}^{start}, j \in [K], s \in [S] \quad (24)$$

$$\sum_s (1 - tl_j) \cdot X_{jts} + I_{j,t-1} - IB_{j,t-1} = d_{jt} + I_{jt} - IB_{jt}, j \in [K], t \in [T] \quad (25)$$

Variables W_{jts}^{start} and W_{jts}^{end} define the respective production windows, which in turn limit the production amounts (X_{jts}) – see requirements (20)–(21). These variables are upper bounded by (22) and lower bounded by (23). Figure 5 provides some examples to explain how these constraints work and to show different production windows definitions. In case 5(a), the slot crosses period t entirely, and therefore the production window bounds in period t match that period's boundaries. In case the slot ends in the middle of time period t , cases 5(c) and (d), the ending date of the production window is the same as the ending time of the slot. Note that the production window is contained within the slot, as it discards the setup time. In case slot s does not intersect period t (i.e. $E_{st} = 0$), requirements (22) and (23) are non-active. The relation between the production window of grade j in slot s and the slot length is established in (24) – the slot contains the production window and the setup times consumed in a grade changeover. Then, the fulfilment of demand is obtained by (25), which allows backlogging orders to be met in subsequent periods.

4.4. Objective function

Given that the system works with an MTO policy and the plant operates continuously on a 24/7 basis (no idle or extra times), orders may backlog at some points in time. The goal of operational production planning and scheduling is then to minimize that backlog at the minimum possible cost. Operational costs include grade setup, production and holding costs. Production costs are mainly related to diluting the costs of the capital intensive equipment, which is achieved by maximizing its utilization (sometimes producing beyond fixed orders). The setup costs, which

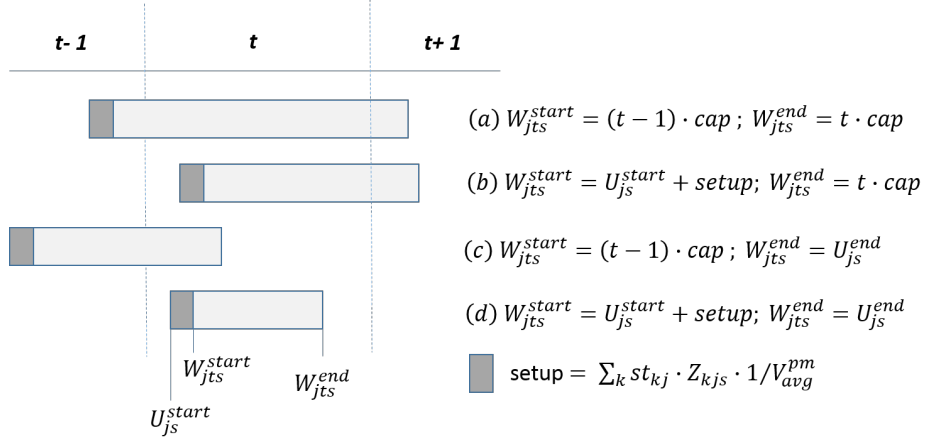


Figure 5: Production window of a slot definition in each time period.

are sequence dependent, concern not only the time and product lost in that process, but also the wear to equipment. Equation (26) represents a cost function that aggregates these main objectives.

Parameters

bc	Cost of backlogging one tonne of paper per period
hc	Cost of holding one tonne of paper in stock per period
sc_{kj}	Setup cost of changing from grade j to k
po	Benefit (negative cost) of producing paper

$$\min \sum_j \sum_t bc \cdot IB_{jt} + \sum_j \sum_t hc \cdot I_{jt} + \sum_j \sum_k \sum_s sc_{kj} \cdot Z_{kjs} - \sum_j \sum_t \sum_s po \cdot X_{jts} \quad (26)$$

Other objectives are considered in the complete model. The driver of the P&P company is not only to produce paper, but also to generate and sell energy based on the steam produced (which is somehow proportional to the throughput of the digester). Therefore, any plan also aims at maximizing the production of steam. Further terms of this objective function are revealed in Subsection 4.6, where some practical issues are addressed.

4.5. Valid inequalities

This subsection introduces sets of valid inequalities to tighten the aforementioned mathematical formulation.

$$\sum_s E_{ts} \geq 1, t \in [T] \quad (27)$$

$$\sum_t E_{ts} \geq 1, s \in [S] \quad (28)$$

$$\sum_s E_{ts} \leq slots_{max}, t \in [T] \setminus \{T\} \quad (29)$$

$$Z_{jjs} \leq 1 - Z_{jku}, j, k \in [K] : k \neq j, s, u \in [S] : u > s \quad (30)$$

$$Z_{jjs} \leq E_{Ts}, j \in [K], s \in [S] \quad (31)$$

Each time period is intersected by at least one slot, as pointed out by (27). Clearly, each slot has to be assigned to at least one time period (28). The maximum number of slots per period is reinforced by (29). Two consecutive non-empty slots produce different grades. In case there is no need to use all the pre-defined slots, the null slots should be placed at the end of the horizon. Requirements (30) and (31) force the fictitious changeovers (Z_{jjs}) to be moved to the end.

4.6. Practical constraints

In order to be used in practice, the model presented above needs to incorporate real-world extensions. Some of the new features are related to customers' satisfaction, while others have to do with managing the resources in the smoothest way possible.

4.6.1. Sales requirements

To consider different types of orders with different priority levels, it is necessary to disaggregate production, inventory and backlog variables. That way, it is possible to introduce different penalties for backlogging specific orders. The suffix *norm* of the new variables refers to the normal priority level, whereas *ship* to ship orders and *prio* to the highest priority orders. Given that multiple ships can be scheduled within the planning horizon, an index $m \in [M]$ is introduced to represent them. For instance, variable X_{jtm}^{ship} denotes the amount of grade j produced in period t to be distributed in ship m . The latest production time of an order to meet the departure of ship m is given by dd_m .

$$I_{j0} = I_{j0}^{norm} + I_{j0}^{prio} + \sum_m I_{jm}^{ship}, j \in [K] \quad (32)$$

$$\sum_s (1 - tl_j) \cdot X_{jts} = X_{jt}^{norm} + X_{jt}^{prio} + \sum_m X_{jtm}^{ship}, j \in [K], t \in [T] \quad (33)$$

$$\sum_{t=1}^{dd_m} X_{jtm}^{ship} + I_{jm}^{ship} - IB_{jm}^{ship,ini} = d_{jm}^{ship} - IB_{jm}^{ship}, j \in [K], m \in [M] \quad (34)$$

$$I_{jT}^{norm} - IB_{jT}^{norm} = d_{j,T+1}^{norm} + I_{j,T+1}^{norm} - IB_{j,T+1}^{norm}, j \in [K] \quad (35)$$

The disaggregation of the initial inventory is performed by (32), where it is allocated to normal, priority or ship orders (the three priority levels described in Section 2). Constraints (33) disaggregate production. Note, however, that there is also aggregation regarding the different slots that cross the period, since at this point we just need to know the completed amounts at the end of each period.

The demand fulfilment equations for normal and priority orders are similar to constraints (25), and for that reasons they are not detailed here. On the other hand, for ships the equations are slightly different – see (34). In fact, each ship m has a single due date (dd_m). Thus, demand fulfilment is not verified in each period, but only in the ship's due date. Therefore, production is summed from the first period until the due date, and no inventory is left after that since there is no other chance to fulfil that demand.

Finally, in (35) the fulfilment of demand beyond the planning horizon is verified. Given that the machine is never idle and paper is not made to stock, if all the demand within the horizon has been met, future demand should be considered. All this demand is then aggregated in an additional term corresponding to period $T + 1$. Naturally, not fulfilling this demand will have a lower impact on the objective function.

4.6.2. Paper campaigns

Here, we describe some constraints related to operational issues in executing paper campaigns. Parameters N_{min} and N_{max} define the lower and upper bounds for the length (in time units) of each grade campaign, respectively. Furthermore, L_{min}^{VLB} refers to the minimum production amount of each VLB-type campaign, $|K_{VLB}|$ to the number of VLB grades, and $dist_{min}$ to the minimum distance between campaigns of the same grade.

$$Z_{jks} = 0, \quad s \in [S], j \in [K], k \in [K_j^{na}] \quad (36)$$

$$N_s \leq N_{max} \cdot \sum_k \sum_{j \neq k} Z_{kjs}, \quad s \in [S] \quad (37)$$

$$N_s \geq N_{min} - N_{max} \cdot \sum_j Z_{jjs}, \quad s \in [S] \quad (38)$$

$$(1 - tl_j) \cdot \sum_t X_{jts} \geq L_{min} \cdot (Y_{js} - Z_{jjs}), \quad j \in [K], s \in [S] \quad (39)$$

$$\sum_{j \in [K_{VLB}]} \sum_t \sum_{s'=s}^{s+|K_{VLB}|-1} X_{jts'} \geq L_{min}^{VLB} \cdot \sum_{k \notin [K_{VLB}]} \sum_{j \in [K_{VLB}]} Z_{kjs}, \quad j \in [K], s \in [S] : s \leq S - |K_{VLB}| + 1 \quad (40)$$

$$\sum_{s'=s}^{s+dist_{min}} (Y_{js'} - Z_{jjs'}) \leq 1 + viol^{dist}, \quad j \in [K], s \in [S] : s \leq S - dist_{min} \quad (41)$$

The first constraints have to do with sequences that cannot be executed. In our case, quality “A” grades should not be produced after basic ones. Indeed, if they do not meet the required quality level, but there is a basic quality campaign following, they can still be used as basic quality products, and the “A” quality has another opportunity to be produced. Therefore, for each grade j , there is a list of forbidden grades $[K_j^{na}]$, which correspond to the better quality versions (if any).

Constraints (37) and (38) define maximum and minimum lengths for campaigns, which should be respected for the sake of the plant's stability. Note that for all fictitious changeovers (Z_{jjs}), the length is zero. Minimums are also defined in terms of the amounts produced (L_{min} and L_{min}^{VLB}) – see (39) and (40). The former applies a minimum lot to individual campaigns, whereas the latter does that to a set of campaigns of the same type (VLB grades).




Figure 6: Example of a plan not following the operational requirements.

Finally, (41) forces a minimum distance ($dist_{min}$) between campaigns of the same grade. The violation of this soft constraint ($viol^{dist}$) is penalized in the objective function. This issue is somehow avoided with setup costs. However, these constraints penalize more explicitly the production of the same grade twice in a short period of time. This type of constraints helps better tailoring the plans to meet industrial requirements, which would otherwise require non-linear cost functions.

Figure 6 provides an example of a plan that does not comply with the operational constraints described here. The corrected version of the plan is that represented in Figure 2. Note that the first part is fixed, since those campaigns were already programmed in the cutting stage. The first campaign, which seems not to comply with the minimum lot size, is indeed in progress. The above constraints are thus only applied after the last fixed campaign.

4.6.3. Production rates

The production rates of all resources should be as smooth as possible to avoid large energy expenditures and wear of the equipment. Here, we illustrate how the digester's rates can be smoothed in the optimization process. Variables V_s^{dig} introduce the speed of the digester in slot s .

$$V_s^{dig} - V_{s-1}^{dig} = \delta_s^{dig+} - \delta_s^{dig-}, s \in [S] \quad (42)$$

$$X_s^{virg} = \alpha \cdot V_s^{dig} \cdot N_s, s \in [S] \quad (43)$$

The rate variations are assessed with equations (42). Minimizing the variables δ_s^{dig+} and δ_s^{dig-} in the objective function will smooth rate variations. The former variable will then assess the rates' increase, whereas the latter will assess their decrease. These equations require an explicit variable for the digester's rate (V_s^{dig}), as opposed to the implicit definition in constraints (8). That results in a non-linear formulation, since the rate variable is multiplied by the slot's length – see (43). Therefore, in order to avoid the complexity of the non-linear formulation, these constraints are only added in a post-optimization step where the slots are fixed.

4.6.4. Recommended tank levels

In practice, the stock should not come close to the physical limits of the tank, since it may trigger operational difficulties and would leave the system more vulnerable to disturbances. For those reasons, recommended levels (below the capacity upper limit and above the lower limit) are established. At some points in time the stock may range between the desirable and physical limits, but the violation of the recommend levels is penalized in the objective function. The stock of virgin pulp in the tank in slot s is now given by $I_s^{virg} + I_s^{virg,rec+} - I_s^{virg,rec-}$, where variables $I_s^{virg,rec+}$ denote the amount of stock above the recommended upper limit, and $I_s^{virg,rec-}$ the amount of stock below the recommended lower limit. In addition, parameter I_{slack-}^{virg} (I_{slack+}^{virg}) refers to the slack between the upper (lower) recommended level and the upper (lower) physical limit.

Equations (9)–(10) should then be replaced by the following:

$$I_{min}^{virg} + I_{slack-}^{virg} \leq I_s^{virg} \leq I_{max}^{virg} - I_{slack+}^{virg}, \quad s \in [S] \quad (44)$$

$$I_{min}^{virg} \leq I_s^{virg} + I_s^{virg,rec+} - I_s^{virg,rec-} \leq I_{max}^{virg}, \quad s \in [S] \quad (45)$$

$$X_s^{virg} + I_{s-1}^{virg} + I_{s-1}^{virg,rec+} - I_{s-1}^{virg,rec-} = O_s^{virg} + I_s^{virg} + I_s^{virg,rec+} - I_s^{virg,rec-}, \quad s \in [S] \quad (46)$$

Equations (44) force the regular variable (I_s^{virg}) to respect the recommended levels, while extra level variables ($I_s^{virg,rec+}$ and $I_s^{virg,rec-}$) are included in the physical limits verification and material balance equations – constraints (45) and (46), respectively.

4.6.5. Fixed schedule

A critical feature of the model, and one of the most important reasons that led to the choice of CGLSP as the basic formulation, is the possibility of fixing production amounts without the need to heuristically define the number of slots required for each campaign. In fact, as slots are able to cross periods, only one slot is required per campaign. Therefore, both the sequence of campaigns and their lot sizes can be easily fixed. The former is a simple assignment ($Y_{js} = 1$, for the respective grade and slot), while the latter is represented in constraints (47). These constraints are essential to freeze campaigns in the DSS that have already been programmed in the cutting stage, i.e. their cutting patterns have been optimized.

$$\sum_t X_{jts} = X_s^{fix}, \quad s \in [S] < prodS, j \in [K] : Y_{js} = 1 \quad (47)$$

X_s^{fix} is the amount to be produced in slot s and $prodS$ is the number of slots with fixed production. This latter parameter has to be taken into account in some of the constraints presented previously.

4.6.6. Stoppages

Finally, in this last subsection we show the way stoppages were modelled. These stoppages can occur in different production resources, such as the digester and the paper machine. We consider all the stoppages' starting and ending times as parameters $stop_r^{start}$ and $stop_r^{end}$, where r is the index of the stoppage. Then, we use binary variables P_{rs}^{start} and P_{rs}^{end} to state if stoppage r starts at the beginning or ends at the end of slot s .

$$stop_r^{start} - hcap \cdot (1 - P_{rs}^{start}) \leq S_s^{start} \leq stop_r^{start} + hcap \cdot (1 - P_{rs}^{start}), \quad r \in [R], s \in [S] \quad (48)$$

$$stop_r^{end} - hcap \cdot (1 - P_{rs}^{end}) \leq S_s^{end} \leq stop_r^{end} + hcap \cdot (1 - P_{rs}^{end}), \quad r \in [R], s \in [S] \quad (49)$$

$$\sum_s P_{rs}^{start} = 1, \quad r \in [R] \quad (50)$$

$$\sum_s P_{rs}^{end} = 1, \quad r \in [R] \quad (51)$$

$$\sum_{s'=1}^s P_{rs'}^{start} \geq \sum_{s'=1}^s P_{rs'}^{end}, \quad r \in [R], s \in [S] \quad (52)$$

Constraints (48) and (49) align the starting and ending times of slots with the starting and ending times of stoppages, respectively. Then, (50) and (51) ensure that only one slot is aligned with the start and/or the end of each stoppage. The last constraints prevent a stoppage from ending before its start.

Variables P_{rs}^{start} and P_{rs}^{end} are then used to constraint the output of the corresponding resources. In the case of the digester, for instance, equations (8) would include these variables to enable or disable the maximum and minimum production rates.

5. Solution method

The solution strategy followed is based on a heuristic solution of the problem's mathematical formulation. The reasons behind this choice come both from practical implications and algorithmic aspects. First and foremost, the large scale and complexity of the model described in the previous section for a regular real-world instance prohibited the use of a commercial solver for the complete model formulation because of its difficulty in even finding a feasible solution in the time limit imposed by the DSS usage. The model suffers from its computational intractability, especially when the number of time slots defined increases. Second, heuristics based on mathematical programming techniques, also known as matheuristics, are a powerful framework to explore the problem's structure and take advantage of today's computational and commercial solvers' power. A survey of these methods is given by Puchinger and Raidl (2005).

Matheuristics trade-off the solution quality obtained by solving the complete model formulation in a commercial solver with the efficiency of heuristics and meta-heuristics. Moreover, these algorithms yield substantial advantages when compared to traditional heuristics, as they require less parameters, and consequently a much lower effort in tuning parameters, and can adapt to changes in the mathematical formulation with limited or zero adjustments. Matheuristics often provide quasi-optimal solutions for a variety of problems. For most companies, having a very good solution is sufficient. Optimality often provides only a marginal improvement for a high additional effort.

Decomposition methods are a popular category of matheuristics. The idea is to decompose the problem into sub-MIPs of manageable size and solve them iteratively. The most popular decomposition approach is based on the time horizon and is the idea of the so-called rolling-horizon algorithms (Dillenberger et al., 1994; Dimitriadis et al., 1997). These methods have been successfully applied in other decision support systems (Guimarães et al., 2014a). Here we apply two decomposition approaches: one based on time slots, and another based on paper grades.

Our matheuristic has three main phases: initial solution, forward-pass and neighbourhood search. All the phases decompose the problem or reduce its dimension by working on the set \mathcal{Y} defined by the variables Y_{js} (setup of grade j in slot s) of the original problem. The overall procedure is illustrated in Figure 7.

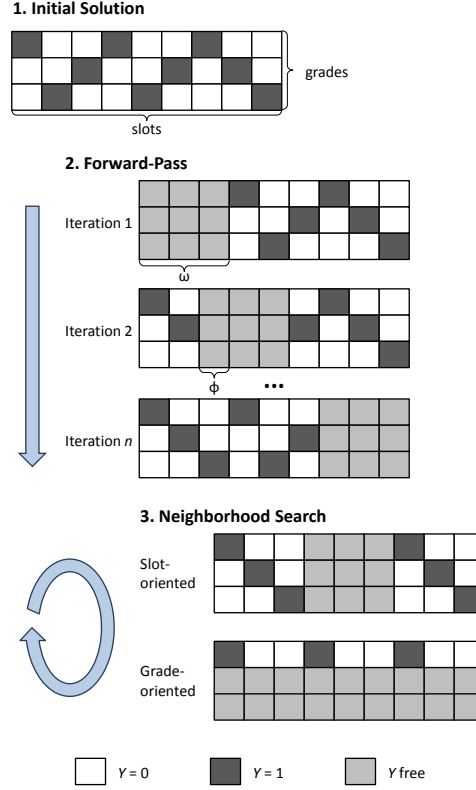


Figure 7: Solution method consisting of three phases: initial solution, forward-pass and neighbourhood search. A sub-MIP is solved in each iteration.

In the first phase, the goal is to generate an initial feasible solution to the problem. For that, the variables in set \mathcal{Y} are fixed in the mathematical formulation, according to an initial sequence provided by a company's manager or a standard (pre-defined) sequence considering the current grade on the paper machine. The value for the remaining variables in the problem is obtained by solving the resulting sub-MIP.

The idea of the second phase is to improve the quality of the initial solution using a forward pass over the continuous time slots. At each iteration of the second phase, a total of ω adjacent slots are re-optimized in a sub-MIP where their corresponding Y_{js} are set as binary, while the remaining Y_{js} variables keep their current value. We start at the beginning of the planning horizon and re-optimize the solution corresponding to the first ω slots. Between iterations the selection of the slots also uses ϕ as the number of overlapping slots, allowing for a smoother schedule. Thus, in the next iteration the solution for the first $\omega - \phi$ slots is fixed and we re-optimize the

flowing ω slots. This process is repeated until the final slot is reached.

The third phase is a neighbourhood search that iteratively decomposes the overall problem. Two neighbourhoods were defined: slot and grade oriented. The slot oriented neighbourhood is similar to the second phase where all the Y_{js} variables associated with a given sub-set of slots are “free” in the model formulation, keeping the other slots unchanged. The main difference lies on the fact that the exploration of slots to re-optimize does not follow a forward pass search, but a probabilistic selection. The grade oriented neighbourhood optimizes the variables Y_{js} of a given sub-set of grades over the entire planning horizon. The choice of the slots and grades is guided by a procedure which uses the frequency and recency of selection to bias the sub-sets. The more frequently or recently a slot or grade is selected, the lower the probability of it being selected to be part of the sub-set to re-optimize. The exploration of a neighbourhood corresponds to the sub-set creation and model re-optimization being performed until a maximum number of consecutive iterations without improvement has been reached. The neighbourhood search alternates between the two neighbourhoods, stopping when both are unable to yield any improvement or the user defined time limit is reached.

6. Decision support system

This section starts by describing the system architecture and how the optimization component is integrated in that framework. Then, the focus is shifted to the system’s usage, where the features and information relevant to managers are detailed. In particular, specific activation and deactivation rules are described, as well as the key performance indicators (KPIs) of the plant.

6.1. System architecture

The DSS was built following a typical setup of full-fledged “Software as a Service” (SaaS) applications. SaaS applications are centrally hosted, and by porting this feature into a service-oriented DSS, we address many problems related to the maintenance and evolution of software. It is especially helpful when phasing out older versions (Gold et al., 2004). Updates and bug fixes are made available in a transparent and immediate way, requiring no further action from the user.

The architecture of this SaaS DSS is comprised of three main layers: the Data Layer, the Optimization Layer and the Web Presentation Layer. The layers communicate and work in a seamless way. Also, the DSS infrastructure interacts with the enterprise resource planning (ERP) and related systems, following the best practices in the field (Harjunoski et al., 2014). Figure 8 gives an overview of the system’s architecture.

The Data Layer component is in control of establishing the necessary connections with third-party systems, namely with the manufacturing execution system (MES) to retrieve information about the shop floor; and with the ERP, which contains typical customer and order management figures. The connection is established over an on-demand VPN link to the data warehouse, after which data is retrieved and mapped to a local data model schema stored in an SQL database (SQL Server).

Data retrieved from the external systems is then made available to the Optimization Layer, providing critical input parameters to the optimization algorithm. Optionally, data can also be provided through an Excel spreadsheet. Excel is still a popular and ubiquitous tool for assembling data from several sources and for assisting in such an iterative and time-consuming process (Harjunoski et al., 2014), despite being potentially more error-prone since users can manipulate

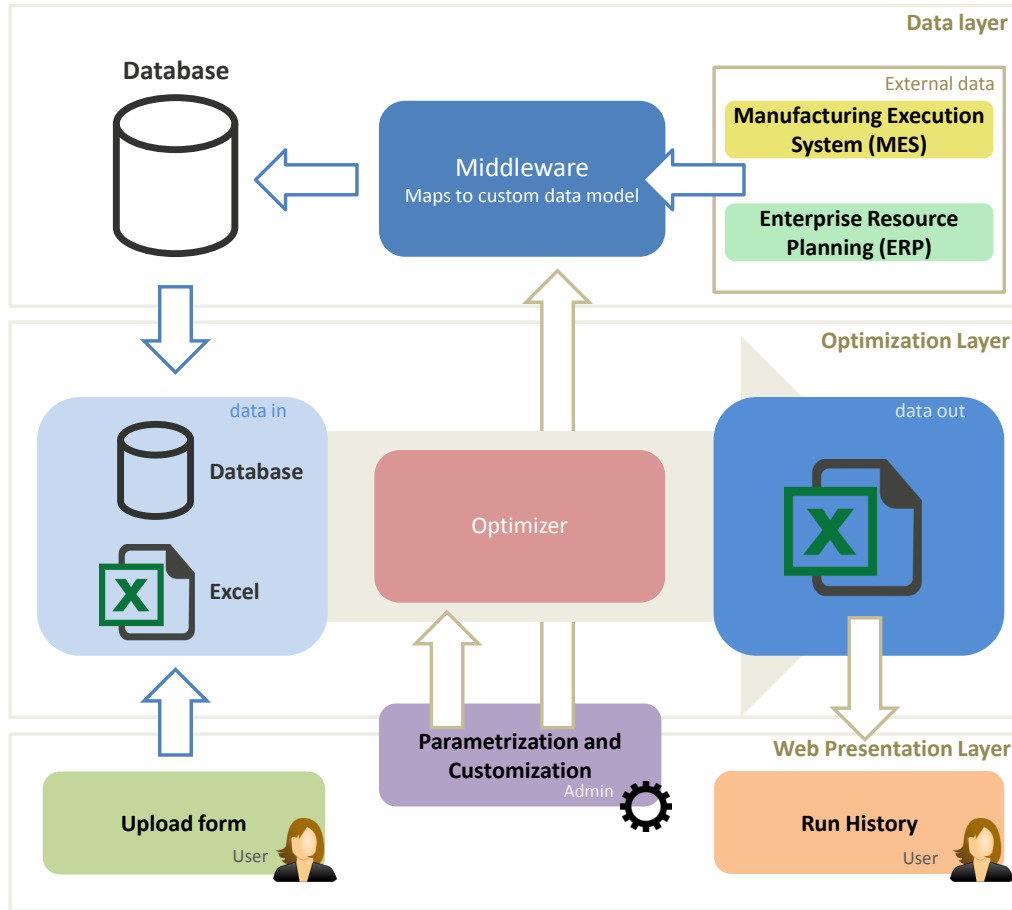


Figure 8: Overview of the DSS architecture, composed by three main layers: Data, Optimization and Web Presentation.

it freely. Regardless of the input source chosen, the data is fed to the Optimizer, which will then generate a solution and provide it as an Excel spreadsheet file.

The optimizer block (illustrated in Figure 9) generates a production plan by executing the optimization algorithm described in the previous section, followed by a post-optimization phase, where the time slots' lengths are fixed and the production rates are smoothed (cf. Subsection 4.6.3). The overall optimization process is wrapped by pre- and post-processing stages. In the former, the size of the mathematical model is reduced by removing paper grades with no demand or backlog to meet. In addition, orders are aggregated by grade and priority level to match the model's parameters. After the optimization, the production amounts have to be disaggregated again, in order to make sure that individual orders are met, and KPIs have to be computed. The allocation of production to orders is achieved by applying the "earliest due date" rule.

None of the layers is directly accessible by the operator of the DSS. Instead, the user interacts with the system through the Web Presentation Layer, where he can upload Excel files to be later

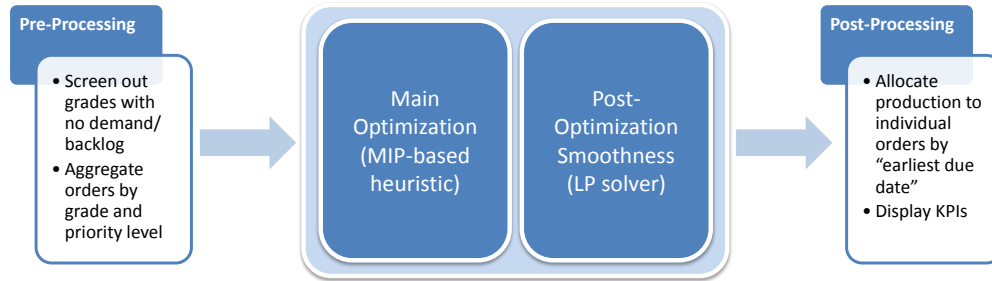


Figure 9: The optimizer component includes the main and post-optimizations, as well as pre- and post-processing stages.

used by the Optimization Layer; consult the run history; download output files; or check the log files. Moreover, advanced users with administrator privileges will also be able to change certain configuration parameters that may interfere directly with the optimizer. In the limit, the web interface could be used only as a gate to the algorithm, a situation in which the user would only upload the Excel input file and then download the Excel output file generated. However, the full potential of the web interface comes with the important information directly retrieved from the ERP systems, along with the tools developed and incorporated in the DSS to assist in the analysis. One of such tools is an advanced table editor (*Handonstable*) that allows in-place editing and direct copy-paste from and to Excel. The main advantage is that critical business logic constraints and error-checking can be implemented at the interface level, preventing potential mistakes from propagating to the Optimization Layer or to be stored in the local database.

6.2. System usage

From a user perspective, our DSS assists both production planners and production managers in obtaining the best production plans for a short-term planning horizon. The process to obtain such plans is presented in Figure 10. As previously mentioned, in order to start the user needs to select which type of data to use — Excel or database.

By choosing to use the database, the user further needs to select a profile (production planner or production manager), and he may also fine-tune some parameters. This parametrization is made directly on the system's interface, which is divided in two main sections. One displays the main steps to generate a plan, and comprises only the frequently changed parameters, in order to expedite the process (Figure 11). The other section contains more structural parameters, such as capacities and processing times, which are more stable and hence accessed only sporadically (Figure 12).

These parameters are organized in different tables, which are interconnected. For instance, when a grade is added to the corresponding table, the table of setups is updated. As illustrated in the figure, this table also presents a gradient of colors which helps managers having a better understanding of the different costs associated with the process.

In the plan generation interface, managers start by customizing the orders (e.g. cancelling, assigning priorities), which are read from the ERP. A variety of filters can be applied to facilitate this characterization. Then, managers check planned production quantities on the spot, as well as their estimated starting and ending times. This table is used to fix the initial part of the plan (essential for rescheduling purposes) and scheduling stoppages (cf. Subsection 4.6.6). To do that,

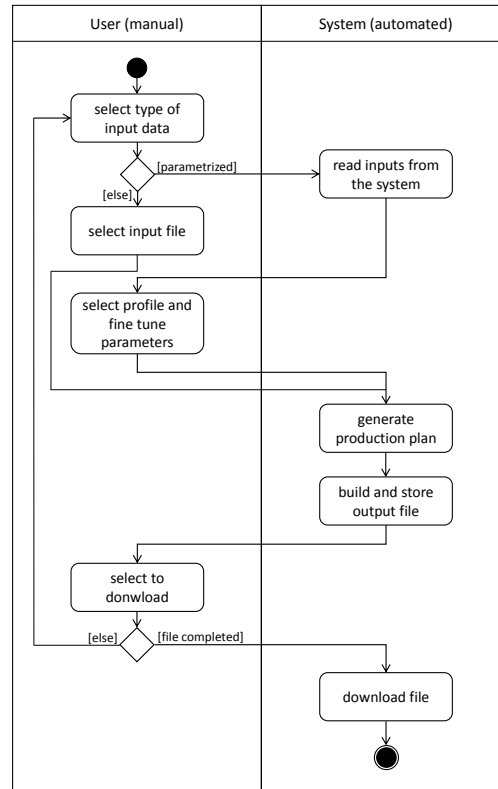





Figure 10: Process of generating a production plan as an activity diagram. The activities on the right-hand side are automated in the decision Support system, and the actions on the left-hand side require an input by the user.

managers can insert new rows at any position of the table and assign either a specific paper grade and production amount or a stoppage (with specified starting and ending times). In addition, a variety of rules can be activated or deactivated, depending on the user's judgement of the current state of the plant. These options include:

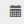
- distinct priority levels (cf. Subsection 4.6.1);
- required quality sequence (cf. Subsection 4.6.2)
- maximum and minimum lot sizes (cf. Subsection 4.6.2);
- maximum and minimum campaigns' lengths (cf. Subsection 4.6.2);
- minimum distance between campaigns (cf. Subsection 4.6.2);
- production rates smoothness (cf. Subsection 4.6.3);
- initial backlog clearance (imposing the fulfilment of the initial backlog orders until the end of the planning horizon).



New Process

Please choose your profile Planner Execute Process

Planning start date 30-10-2014 15:55 

[Orders](#)
[Ships](#)
[Plan](#)
[Other settings](#)
[Uniformance](#)

Paper Machine

Grades	Production (Tons)	Start time (estimated)	End time (estimated)	Optivision
KLB225	47.44	30-10-2014 15:55	30-10-2014 16:53	✓
KLB200A	56.78	30-10-2014 16:53	30-10-2014 18:16	✓
KLB200	434	30-10-2014 18:16	31-10-2014 03:11	✓
KLB170	851.91	31-10-2014 03:11	31-10-2014 20:56	✓
KLB125	324.12	31-10-2014 20:56	01-11-2014 05:03	✓
KLB115	401.61	01-11-2014 05:03	01-11-2014 15:50	✓
VLB165	436.65	01-11-2014 15:50	02-11-2014 01:22	✓
VLB115	460.53	02-11-2014 01:22	02-11-2014 13:36	✓
VLB140	180.42	02-11-2014 13:36	02-11-2014 17:33	✓

☒ Fixed sequence

Last update to Optimision 17-09-2014 00:22:20 Update

Figure 11: Web interface of the DSS for generating a production plan (only frequently changed parameters are displayed).

When the setup of the input data is concluded, the system starts to generate the production plan by running the Optimizer, which includes the MIP-based heuristic (cf. Section 5). After running this module, the system builds and stores the output excel file that will serve as the basis for the decision process. Figure 13 is an example of an output. The output file comprises different spreadsheets for the multiple areas of the P&P mill, represented by different colors in the figure. The first sheets concern the entire mill and display all the computed production figures in tables, as well as relevant KPIs, such as average and final backlog of orders, equipment utilization ratios and total production output. The area-specific sheets then represent Gantt charts depicting the sequence of operations, trend charts displaying the forecasted evolution of inventories of intermediate and final products, and a comprehensive reporting on order fulfilment with expected dates of delivery. The fact that all charts are generated from the tables and the latter contain formulas for a great part of their figures (for instance, intermediate inventories are not fixed values in the spreadsheet, but depend on the respective production and consumption values) allows a convenient adjustment of the resulting schedule to address unsolved business/production issues or to perform *what-if* analysis. At the end, if the user wants to perform some sensitivity analysis he may run another production plan or download the output file to his computer.

The system's usage described can be performed in three main contexts:

- periodic planning – the managers typically determine an optimized plan on a weekly basis

and revise it daily;

- reactive scheduling – in the presence of a disturbance situation which causes a serious disruption to the original plan, recover from that disturbance using an appropriate sequence of campaigns (sometimes not obvious to plant managers);
- *what-if*-analysis – supporting planning in hierarchically superior levels by checking the effectiveness of practical rules, expanding storage and production capacities (and other retrofit designs), changing weights of different objectives and scheduling maintenance.

The second context tackles both internal uncertainties (such as equipment disturbances) and external uncertainties (such as product demand) in a reactive way. Proactive measures can also be applied, but they have to be defined manually. Thus, backlog situations may be prevented by increasing the demand considered (there is a parameter for that in the planning options tab), and disturbance propagation can be mitigated by narrowing the recommended levels of the intermediate tanks (cf. Subsection 4.6.4).

7. Results

In this section we start by looking at the usage of the system during a period of four months, and then delve into one plan generated by the system and its differences comparatively to the manually generated counterpart. In the end, a plan considering a long stoppage is analysed.

7.1. Overall results of the system's usage

Production planners and managers are not willing to wait more than 15 to 20 minutes for an optimized plan in normal conditions. When facing an operational issue, such as a disturbance, that time can be reduced to 5 to 10 minutes. With that time limit, state-of-the-art solvers such as Cplex are not even able to find a solution for a regular instance. The focus here is thus much more on feasibility and improvement over current practice than on proving optimality. For that reason, a considerable amount of constraints are relaxed (becoming soft constraints) so that our constructive heuristic is always able to find an initial solution. Then, the MIP-based heuristic improves that solution to a point that satisfies plant managers, that is where soft constraints are satisfied as much as possible and the plan performs well according to the objectives defined.

Our system was made available online for the company to use. Various iterations were performed with the users in order to fine tune all the parameters and introduce additional constraints to the mathematical model, so that the resulting plans met their requirements. Table 1 contains information about 22 runs executed during four months after those iterations. Even though the total number of grades is 29, the number of grades with demand or backlog to satisfy at each moment is reduced to 16/17. This shows the importance of the pre-processing steps described in the previous section. The maximum number of campaigns (i.e. time slots) is set to 35 (for a horizon of 15 days) and the model can use them all or leave some empty. The fixed (frozen) campaigns depend on what was already scheduled in the cutting stage and on additional considerations managers might want to make.

The table reports the improvements achieved over the initial solution in runs executed by the planners. We can see that those values are highly variable. Essentially, they will depend on whether the initial solution is already good or not, and on the time limit given, which is defined by the user in every run. Nevertheless, an average improvement of 34% was obtained, with

Table 1: Company's usage of the system from June to September 2014.

#	Date	Grades	Fixed campaigns	Total campaigns	Improv over ini sol	Time (min)
1	09/06/2014	17	4	35	75%	13
2	17/06/2014	17	7	34	81%	12
3	19/06/2014	17	9	28	26%	8
4	20/06/2014	17	8	33	51%	10
5	20/06/2014	17	20	27	5%	4
6	27/06/2014	17	10	30	29%	10
7	04/07/2014	17	14	28	0%	7
8	09/07/2014	16	7	35	4%	10
9	14/07/2014	16	3	31	1%	18
10	18/07/2014	17	4	35	45%	12
11	18/07/2014	17	27	28	3%	2
12	22/07/2014	17	9	32	8%	11
13	24/07/2014	17	4	30	6%	12
14	24/07/2014	17	7	31	7%	15
15	24/07/2014	17	13	29	63%	9
16	29/07/2014	17	5	30	87%	18
17	06/08/2014	17	5	30	93%	14
18	06/08/2014	16	5	32	46%	16
19	27/08/2014	17	9	33	55%	11
20	03/09/2014	17	7	29	6%	10
21	15/09/2014	16	6	31	18%	15
22	25/09/2014	17	9	29	48%	14
Avg	–	17	9	31	34%	11

an average running time of 11 minutes, which shows the ability of the heuristic to correct and improve production plans in short periods of time.

To validate the quality of those plans, they need to be compared to those generated manually by the company. The system can also be useful in this regard. In fact, the system can be used not only to optimize a plan from scratch, but also to test certain schedules by fixing the sequence and amounts of paper campaigns. The table clearly shows that the company has applied these two usages, comparing the output of both (for instance, runs 4 and 5, which have very different numbers of fixed campaigns – run 4 refers to the optimum plan and run 5 to the so-called manual solution). When fixing most campaigns, the algorithm is naturally much faster to finish. However, that should not be a reason for not allowing the system to optimize the plans, since just providing the initial sequence guarantees a good solution in a few minutes, in case the sequence is good, and makes it possible to improve on that solution.

The plan of run 4 (taking full advantage of the DSS) outperforms that of run 5 (manual plan) by 59%. Regarding runs 10 (optimized) and 11 (manual), there was a 71% improvement. These values are unreasonably high since some soft constraints were violated in the company's plans, whereas our system has addressed them. A more insightful comparison is performed in the next subsection, where we look at the schedules and the main KPIs.

7.2. Optimized plan vs. manual plan – an example

Two production plans, one optimized by the system and the other defined manually (although still generated by the system), are now examined. Figures 14 and 15 depict these two plans, in particular the schedule of campaigns and the virgin pulp stock level (measured in tonnes). As usual, the initial part of the plans is fixed. The differences thus start in the middle, where the

optimized plan includes KLB275A and KLB170A (which were ignored by the manual plan) and an additional campaign of KLB225, on the descending part of the sequence. Note that plans tend to follow this smooth ascending and descending pattern of grades, in order to avoid costly setups. Then, the cycle of the virgin pulp level is extended in the optimized plan to produce a larger amount of KLB170 and to avoid another campaign of this grade. After that, the VLB campaigns have to start in order to raise the level of virgin pulp again. In fact, these grades incorporate considerably more recycled fibres and hence are used to balance the KLB grades, which have the opposite composition - more virgin fibres. The manual plan ends with low KLB grades, which were not repeated in the optimized plan.

To better understand the different planning approaches, we need to look at their impact on the main KPIs. The relative improvement of the optimized plan over the manual counterpart is given in Table 2.

Table 2: Comparison of the manual and optimized plans, with respect to main KPIs.

Backlog	Setups	Production	Stock	Overall
-6.4%	-10.0%	-1.8%	-2.6%	-7.4%

From the table, one may conclude that the additional campaigns in the optimized plan helped to improve backlog by more than 6%. Moreover, the smoother grade changeovers and the absence of the second wave of KLB grades at the end of the manual plan (where KLB135 restarts) made it possible to reduce setup costs by 10%. Furthermore, inventory was reduced by almost 3%. The only advantage of the manual plan is a 2% additional production. The optimization method made a trade-off when extending the duration of the production cycle, which forced a slight reduction of the paper machine's rate (as it was exhausting the pulp tanks), but allowed for a better fulfilment of demand and improvement of operational costs. In manual planning it is difficult to grasp this kind of trade-offs and practitioners are forced to use simple rules, such as a given production cycle, to generate feasible plans. Nevertheless, when programming stoppages for equipment maintenance or when stoppages arise in disturbance situations, it is more difficult to devise rules. In those cases, the DSS can bring additional value.

7.3. Plan considering stoppages

To conclude this section, a plan considering a long stoppage in the entire plant is studied. The plan is illustrated in Figure 16, where the stoppage of the paper machine is represented as another paper grade. The paper machine stops after the digester in order to consume all the pulp left in the tanks. At the start-up, the digester is activated before the machine to create some inventory of pulp. The machine starts with VLB165, which consumes significantly more recycled than virgin pulp. This way, the virgin pulp inventory continues to rise until it reaches a level where the KLB products (which consume more virgin pulp) can be produced.

Planning a stoppage takes the resources to their actual limits. Therefore, the rates have to be carefully managed. When generating the illustrated plan, some simplifying assumptions were made, such as the ability of the production rates to rise from (or fall to) zero instantly. The aim there was just to obtain a rough idea of what the plan would be. Nevertheless, the system allows considering maximum rates variations, as well as easily redefining maximum and minimum rates. Hence, after this initial plan managers could generate a more accurate plan by further customizing the system inputs.

8. Conclusions and future work

The DSS reported in this paper has been developed to help a P&P company deal with the various challenges of their production planning and scheduling activities. Companies in this sector have to differentiate themselves by providing the best customer service at minimum cost, since paper prices are determined by the market. Therefore, multiple criteria have to be considered and properly weighted when devising a production plan. However, managers do not have the time at the operational level to thoroughly evaluate plans, which must satisfy a multitude of complex constraints. Hence, the DSS supports the automatic generation of plans with different levels of customization (e.g. fixing the sequence and/or amounts of paper campaigns).

Several iterations were performed with the company in order to include a set of essential features for their planning tasks and to design user-friendly interfaces. The immediate updates (due to the SaaS setup), in-place editing, direct copy-paste from and to Excel, input validation procedures, automatic connection to existing systems, and separation of plan generation and fine tuning interfaces are examples of technical details that were crucial for the usability of the system. On the other hand, the optimization model provided the required flexibility in incorporating operational constraints, such as fixing campaigns and scheduling stoppages. The method then makes use of the model to generate good quality solutions in the short periods of time available at the operational level.

As future work, additional features can be included, such as improving the interactivity of the final schedules or providing more customization of the input data. These enhancements should be done according to the company's actual needs. In addition, the modular design of the system, as well as the general-purpose of the optimization method, would make it easier to adapt or extend to other similar plants. It would also be interesting to extend the comparison to a large set of manual plans and perform some sensitivity analysis (regarding the savings obtained by the DSS) on certain model parameters, such as product demands.

In terms of the DSS engine, the formulation may still be tightened by additional valid inequalities or reformulations. Comparing the formulations of Karimi and McDonald (1997) and Camargo et al. (2012) could give some insights in this regard. Moreover, the "facility plant location" reformulation could also improve the model's solvability (Amorim et al., 2011). Another important aspect to be examined is the impact of the soft constraints on the optimization performance, and how that could be addressed by the heuristic method. The latter could also be improved with respect to the neighbourhoods used, possibly allowing it to add/remove campaigns to/from certain positions, as suggested by Figueira et al. (2013).

Finally, as computational power increases and the efficiency of the methods improve, the production planning problem can integrate more details of other activities (up or downstream in the supply chain) or address important issues, such as process variability and disturbances. The latter may have an important impact on the production system, and therefore its consideration in a proactive planning approach is highly relevant.

Acknowledgements

This work is partially financed by the Project NORTE-07-0124-FEDER-000057, supported by the North Portugal Regional Operational Programme (ON.2 – O Novo Norte) under the National Strategic Reference Framework (NSRF). The first author is grateful to the Portuguese Foundation for Science and Technology (FCT) for awarding him the Ph.D. Grant SFRH/BD/80109/2011.

Appendix A. List of acronyms

Industry:

KLB	KraftLiner Board (with more virgin fibres)
MTO	Make-To-Order
P&P	Pulp and Paper
RPM	Revolutions Per Minute
VLB	KraftLiner Board (with more recycled fibres)

Planning and scheduling:

CGLSP	Crossed Grid Lot-sizing and Scheduling Problem
CLSD	Capacitated Lot-Sizing with Sequence Dependent setups
CSLP	Continuous Setup Lot-sizing Problem
DLSP	Discrete Lot-sizing and Scheduling Problem
GLSP	General Lot-sizing and Scheduling Problem
PLSP	Proportional Lot-sizing and Scheduling Problem

Information technology:

DSS	Decision Support System
ERP	Enterprise Resource Planning
KPI	Key Performance Indicator
MES	Manufacturing Execution System
SaaS	Software as a Service
SQL	Structured Query Language
VPN	Virtual Private Network

Analytics:

LP	Linear Programming
MIP	Mixed Integer Programming
OR	Operations Research
PSE	Process Systems Engineering

Appendix B. Mathematical notation

Indices

t	Index for time period ($t \in [T] = \{1, \dots, T\}$)
s, u	Indices for time slot ($s, u \in [S] = \{1, \dots, S\}$)
j, k	Indices for paper grades ($j, k \in [K] = \{1, \dots, K\}$)

Parameters

cap	Period's capacity (in hours)
$hcap$	Planning horizon's capacity (in hours): $hcap = T \cdot cap$
$slots_{max}$	Maximum number of slots crossing one period
$V_{min}^{dig} (V_{max}^{dig})$	Minimum (maximum) rate of the digester (in rpm)
α	conversion factor of digester's throughput (from rpm to tonnes per hour)
$I_{min}^{virg} (I_{max}^{virg})$	Minimum (maximum) level of virgin pulp stocked in the tanks (in tonnes)
f_j	Percentage of water in paper grade j
tlb_j	Average trim loss (in percentage) of paper grade j that goes to the loss pulp tank
b_j^{virg}	Percentage of virgin pulp used in the production of paper grade j
b_{loss}^{virg}	Percentage of virgin pulp in the loss pulp
sl_{kj}	Paper lost in a changeover from grade k to j (in tonnes)
st_{kj}	Time lost in a changeover from grade k to j (in tonnes)
p_j^{min}	minimum time (at machine's maximum rate) to produce one tonne of paper grade j (in hours)
p_j^{max}	maximum time (at machine's minimum rate) to produce one tonne of paper grade j (in hours)
d_{jt}	Demand for paper grade j in period t (in tonnes)
tl_j	Average trim loss (in percentage) of paper grade j
I_{j0}	Initial inventory of paper grade j (in tonnes)
IB_{j0}	Initial backlog of paper grade j (in tonnes)
bc	Cost of backlogging one tonne of paper per period
hc	Cost of holding one tonne of paper in stock per period
sc_{kj}	Setup cost of changing from grade j to k
po	Benefit (negative cost) of producing paper

Decision variables

S_s^{start}	Starting time of slot s
S_s^{end}	Ending time of slot s
U_{js}^{start}	Starting time of slot s for grade j
U_{js}^{end}	Ending time of slot s for grade j
E_{ts}	$\begin{cases} 1 & \text{if slot } s \text{ intersects period } t \\ 0 & \text{otherwise} \end{cases}$
Y_{js}	$\begin{cases} 1 & \text{if the paper machine is set up for grade } j \text{ in slot } s \\ 0 & \text{otherwise} \end{cases}$
N_s	Length of slot s (in hours)
X_s^{virg}	Production of virgin pulp in slot s (in tonnes)
O_s^{virg}	Amount of virgin pulp fed into the paper machine in slot s (in tonnes)
I_s^{virg}	Inventory of virgin pulp at the end of slot s (in tonnes)
X_s^{loss}	Production of loss pulp in slot s (in tonnes)
O_s^{loss}	Amount of loss pulp used in slot s (in tonnes)
X_{jts}	Production of paper grade j in period t and slot s (in tonnes)
Z_{kjs}	$\begin{cases} 1 & \text{if a changeover from grade } k \text{ to } j \text{ occurs at the beginning of slot } s \\ 0 & \text{otherwise} \end{cases}$
W_{jts}^{start}	Starting time of slot s for grade j in period t
W_{jts}^{end}	Ending time of slot s for grade j in period t
I_{jt}	Inventory of paper grade j at the end of period t (in tonnes)
IB_{jt}	Backlog of paper grade j at the end of period t (in tonnes)

References

- Akkiraju, R., Keskinocak, P., Murthy, S., Wu, F., 2001. An agent-based approach for scheduling multiple machines. *Applied Intelligence* 14 (2), 135–144.
- Almada-Lobo, B., Klabjan, D., Carravilla, M. A., Oliveira, J. F., 2010. Multiple machine continuous setup lotsizing with sequence-dependent setups. *Computational Optimization and Applications* 47 (3), 529–552.
- Amorim, P., Antunes, C. H., Almada-Lobo, B., 2011. Multi-objective lot-sizing and scheduling dealing with perishability issues. *Industrial & Engineering Chemistry Research* 50 (6), 3371–3381.
- Amorim, P., Pinto-Varela, T., Almada-Lobo, B., Barbósa-Póvoa, A., 2013. Comparing models for lot-sizing and scheduling of single-stage continuous processes: Operations research and process systems engineering approaches. *Computers and Chemical Engineering* 52, 177–192.
- Bongers, P. M., Bakker, B., 2006. Application of multi-stage scheduling. *Computer Aided Chemical Engineering* 21, 1917–1922.
- Camargo, V., Toledo, F., Almada-Lobo, B., 2012. Three time-based scale formulations for the two-stage lot sizing and scheduling in process industries. *Journal of the Operational Research Society* 63, 1613–1630.
- Castro, P. M., Harjunkoski, I., Grossmann, I. E., 2009. New continuous-time scheduling formulation for continuous plants under variable electricity cost. *Industrial & engineering chemistry research* 48 (14), 6701–6714.
- CEPI, June 2013. Key statistics – european pulp and paper industry 2012. Tech. rep.
- Clark, A. R., Clark, S. J., 2000. Rolling-horizon lot-sizing when set-up times are sequence-dependent. *International Journal of Production Research* 38 (10), 2287–2307.
- Correia, M., Oliveira, J., Ferreira, J., 2012. Integrated resolution of assignment, sequencing and cutting problems in paper production planning. *International Journal of Production Research* 50 (18), 5195–5212.
- Dillenberger, C., Escudero, L. F., Wollensak, A., Zhang, W., 1994. On practical resource allocation for production planning and scheduling with period overlapping setups. *European Journal of Operational Research* 75 (2), 275–286.
- Dimitriadis, A., Shah, N., Pantelides, C., 1997. Rtn-based rolling horizon algorithms for medium term scheduling of multipurpose plants. *Computers & Chemical Engineering* 21, S1061–S1066.

- Drexl, A., Haase, K., 1995. Proportional lotsizing and scheduling. *International Journal of Production Economics* 40 (1), 73–87.
- Erdirik-Dogan, M., Grossmann, I. E., 2008. Simultaneous planning and scheduling of single-stage multi-product continuous plants with parallel lines. *Computers & Chemical Engineering* 32 (11), 2664–2683.
- Figueira, G., Oliveira Santos, M., Almada-Lobo, B., 2013. A hybrid vns approach for the short-term production planning and scheduling: A case study in the pulp and paper industry. *Computers & Operations Research* 40 (7), 1804–1818.
- Fleischmann, B., 1994. The discrete lot-sizing and scheduling problem with sequence-dependent setup costs. *European Journal of Operational Research* 75 (2), 395–404.
- Fleischmann, B., Meyr, H., 1997. The general lotsizing and scheduling problem. *OR Spectrum* 19, 11–21.
- Furlan, M., Almada-Lobo, B., Santos, M., Morabito, R., 2015. Unequal individual genetic algorithm with intelligent diversification for the lot-scheduling problem in integrated mills using multiple-paper machines. *Computers & Operations Research* 59, 33–50.
- Gold, N., Mohan, A., Knight, C., Munro, M., 2004. Understanding service-oriented software. *Software, IEEE* 21 (2), 71–77.
- Guimarães, L., Amorim, P., Sperandio, F., Moreira, F., Almada-Lobo, B., 2014a. Annual distribution budget in the beverage industry: A case study. *Interfaces* 44 (6), 605–626.
- Guimarães, L., Klabjan, D., Almada-Lobo, B., 2014b. Modeling lotsizing and scheduling problems with sequence dependent setups. *European Journal of Operational Research* 239 (3), 644–662.
- Harjunkoski, I., Maravelias, C., Bongers, P., Castro, P., Engell, S., Grossmann, I., Hooker, J., Méndez, C., Sand, G., Wassick, J., 2014. Scope for industrial applications of production scheduling models and solution methods. *Computers & Chemical Engineering* 62, 161–193.
- Janak, S. L., Floudas, C. A., Kallrath, J., Vormbrock, N., 2006. Production scheduling of a large-scale industrial batch plant. i. short-term and medium-term scheduling. *Industrial & engineering chemistry research* 45 (25), 8234–8252.
- Jankunaite, D., 2006. Pulp and paper industry - case study. *SCA Forest Products AB*, <http://www.balticuniv.uu.se/index.php>.
- Karimi, I. A., McDonald, C. M., 1997. Planning and scheduling of parallel semicontinuous processes. 2. short-term scheduling. *Industrial & Engineering Chemistry Research* 36 (7), 2701–2714.
- Karmarkar, U. S., Schrage, L., 1985. The deterministic dynamic product cycling problem. *Operations Research* 33 (2), 326–345.
- Kopanos, G. M., Puigjaner, L., Georgiadis, M. C., 2011a. Resource-constrained production planning in semicontinuous food industries. *Computers & Chemical Engineering* 35 (12), 2929–2944.
- Kopanos, G. M., Puigjaner, L., Georgiadis, M. C., 2012. Efficient mathematical frameworks for detailed production scheduling in food processing industries. *Computers & Chemical Engineering* 42, 206–216.
- Kopanos, G. M., Puigjaner, L., Maravelias, C. T., 2011b. Production planning and scheduling of parallel continuous processes with product families. *Industrial & Engineering Chemistry Research* 50 (3), 1369–1378.
- Lim, M.-F., Karimi, I., 2003. Resource-constrained scheduling of parallel production lines using asynchronous slots. *Industrial & engineering chemistry research* 42 (26), 6832–6842.
- Maravelias, C. T., 2005. Mixed-time representation for state-task network models. *Industrial & engineering chemistry research* 44 (24), 9129–9145.
- Martel, A., Rizk, N., D'Amours, S., Bouchriha, H., 2005. Synchronized production-distribution planning in the pulp and paper industry. In: Langevin, A., Riopel, D. (Eds.), *Logistics Systems: Design and Optimization*. Springer US, pp. 323–350.
- Murthy, S., Akkiraju, R., Goodwin, R., Keskinocak, P., Rachlin, J., Wu, F., Yeh, J., Fuhrer, R., Kumaran, S., Aggarwal, A., Sturzenbecker, M., Jayaraman, R., Daigle, R., 1999. Cooperative multiobjective decision support for the paper industry. *Interfaces* 29, 5–30.
- Poltroniere, S., Poldi, K., Toledo, F., Arenales, M., 2008. A coupling cutting stock-lot sizing problem in the paper industry. *Annals of Operations Research* 157, 91–104.
- Puchinger, J., Raidl, G., 2005. Combining metaheuristics and exact algorithms in combinatorial optimization: A survey and classification 3562, 113–124.
- Respício, A., Captivo, M., 2008. Marketing-production interface through an integrated dss. *Journal of Decision Systems* 17 (1), 119–132.
- Rizk, N., Martel, A., D'Amours, S., 2008. Synchronized production-distribution planning in a single-plant multi-destination network. *Journal of the Operational Research Society* 59 (1), 90–104.
- Santos, M. O., Almada-Lobo, B., 2012. Integrated pulp and paper mill planning and scheduling. *Computers & Industrial Engineering* 63 (1), 1–12.
- Shaik, M. A., Floudas, C. A., Kallrath, J., Pitz, H.-J., 2009. Production scheduling of a large-scale industrial continuous plant: Short-term and medium-term scheduling. *Computers & Chemical Engineering* 33 (3), 670–686.
- Szabó, L., Soria, A., Forsström, J., Keränen, J., Hytönen, E., 2009. A world model of the pulp and paper industry: Demand, energy consumption and emission scenarios to 2030. *Environmental Science & Policy* 12 (3), 257–269.

- van Elzakker, M. A., Zondervan, E., Raikar, N. B., Grossmann, I. E., Bongers, P. M., 2012. Scheduling in the fmccg industry: An industrial case study. *Industrial & Engineering Chemistry Research* 51 (22), 7800–7815.
- Wassick, J. M., 2009. Enterprise-wide optimization in an integrated chemical complex. *Computers & Chemical Engineering* 33 (12), 1950–1963.
- Wassick, J. M., Ferrio, J., 2011. Extending the resource task network for industrial applications. *Computers & chemical engineering* 35 (10), 2124–2140.
- Westerlund, J., Hästbacka, M., Forsell, S., Westerlund, T., 2007. Mixed-time mixed-integer linear programming scheduling model. *Industrial & engineering chemistry research* 46 (9), 2781–2796.



Figure 12: Web interface of the DSS for fine-tuning structural (less frequently changed) parameters.

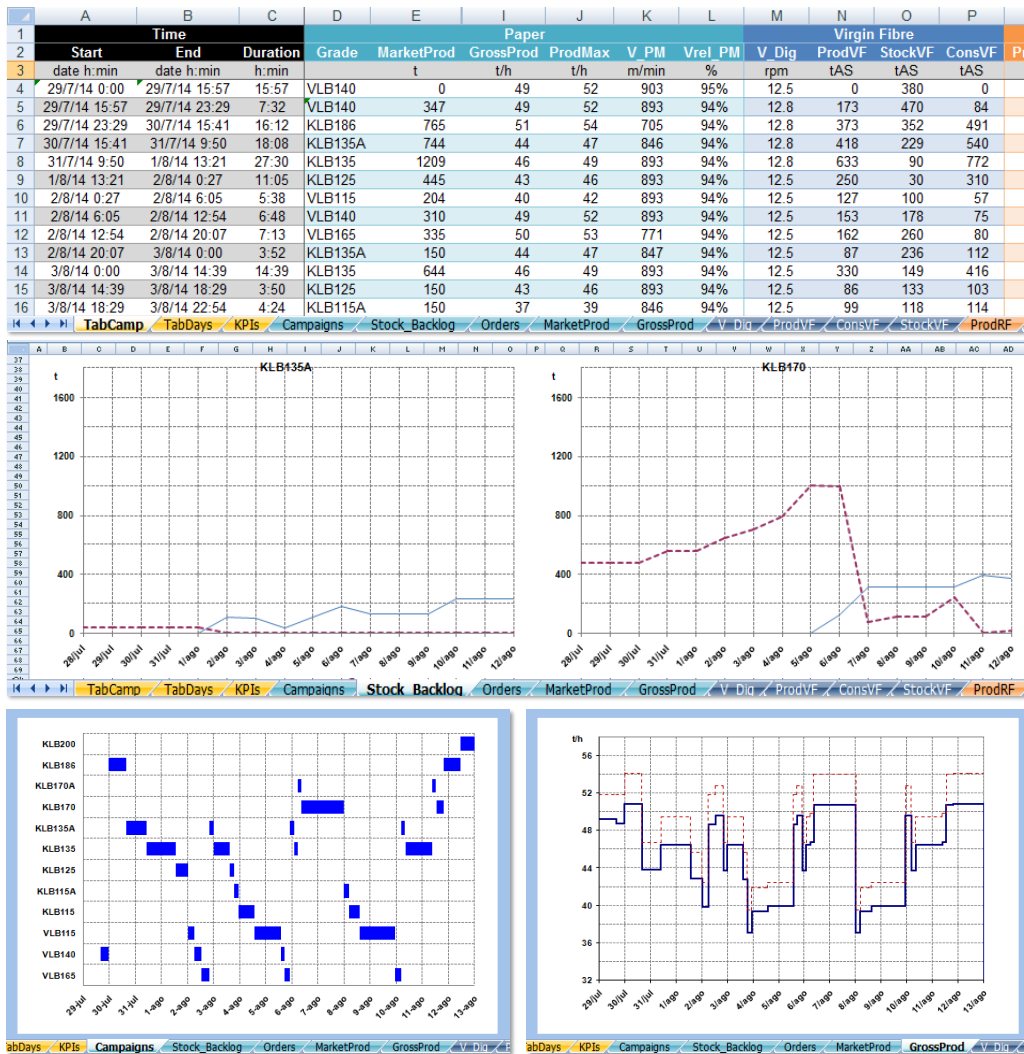


Figure 13: Output Excel file with tables and charts detailing the production plan (complete table, evolution of final products stock / backorders, paper campaigns and machine's production rate, respectively).




Figure 14: Manual plan – generated by the system, fixing the sequence and size of paper campaigns (grades on the left and virgin pulp on the right axis).




Figure 15: Optimized plan – generated by the system, optimizing the sequence and size of paper campaigns (grades on the left and virgin pulp on the right axis).




Figure 16: Plan considering a long long stoppage – paper grades (bars) and digester's rate (dashed line) represented on the left axis and virgin pulp tank (solid line) on the right axis.