

Supply Chain Optimization in Pulp Distribution using a Rolling Horizon Solution Approach

David Bredström^a and Mikael Rönnqvist^b

^aLinköping University, Linköping, Sweden

^bNorwegian School of Economics and Business Administration, Bergen, Norway

Abstract

In this paper we consider a combined supply chain and ship routing problem for a large pulp producer in Scandinavia. The problem concerns the distribution of pulp to customers, with route scheduling of ships as a central part of modeling. It is an operative planning problem with daily ship routing decisions over a 40 days period. The pulp supply is determined by fixed production plans, and the transport flows and storages are modeled with the requirement to satisfy the demand in a cost-optimal way. We develop a mixed integer programming model with binary variables for route usage of a vessel.

The problem is solved with a heuristic solution method, based on a rolling time horizon and a standard branch and bound algorithm. We apply the heuristic on problem instances with real world data, and compare results from reduced problem instances with the results from an exact branch and bound search. The computational experiments indicate that real world problems are solvable with the solution method and that it in many cases can be very efficient.

Keywords: Supply chain, Ships, Scheduling, Mixed integer programming

Introduction

In this paper we consider planning of the supply chain for Södra Cell AB, one of the world's leading manufacturer of market pulp intended for paper production. With five pulp mills in use (three in Sweden and two in Norway), Södra Cell produces more than two million tonnes of pulp per year comprising around 30 different products that are distinguished by brightness and other properties. Södra Cell AB is a subsidiary of Södra, an economic association owned by more than 34,000 forest owners in southern Sweden. Södra's complete supply chain includes harvesting, transportation of wood, production at the mills, storage (at mills, in forests and at harbours), distribution to terminals by vessels, storage at terminals and distribution both directly from mills and from harbours by truck and train to customers located mainly in Europe.

We consider the distribution of pulp to export customers (as well as domestic customers) with route scheduling for vessels as a central part of modeling. The pulp supply is defined by production plans, and the basic model idea is to satisfy the demand in a cost optimal way. Södra Cell delivers large quantities of pulp every year and the distribution is, if measured with costs, one of the major factors of the supply chain. The production and distribution planning is today performed sequentially, beginning with rough production plans for forecasted demand over a time horizon of one year. Given the approximative production plans, an exact schedule for each mill is determined for the upcoming three months and followed by planning the distribution, including ship scheduling for around one month at the time. Södra Cell wants to develop a support system that, given a production plan, can find routes for cargo vessels, can decide when it is beneficial to lease extra vessels on the spot market and determine the pulp quantities to be transported by different means of transportation.

The use of supply chain models within businesses is a rapidly growing area of interest. Shapiro's book¹ provides a good overview of work in this area, while Drexel and Kimms² provide a good survey paper. Very little work has been done on ship scheduling and routing with focus on the whole supply chain.³ The classical papers by Appelgren^{4,5} that cover optimization methods for ship scheduling are among the first papers to appear in the area of ship scheduling. More recent work can be found in the papers of Fagerholt⁶ and Fagerholt and Christiansen.⁷ Corresponding decision support with focus on the forest industry can be found in, for example Carlsson and Rönnqvist⁸ and Bredström et al.^{9,10} Related problems have been studied as the Inventory Routing Problems, IRP. Although the formulations as IRP are similar to the formulation we use in this paper, there are some major differences. First, there are usually only consuming or producing facilities, that is no intermediate storages are considered. Second, the distribution often concerns only one product, while the case of Södra Cell comprises around fifteen products. Third, in our application we may have multiple pickups and multiple deliveries which makes the problem even more general and harder. Examples of IRP solutions can be found in distribution of industrial gases,¹¹ refuse collection,^{12,13} distribution of ammonia¹⁴ and in the car industry.¹⁵

The idea to use a rolling horizon approach on supply chain optimization and planning problems has been applied in different contexts. However, as far as we know it has not been applied for ship scheduling and our type of application. Even though a rolling time horizon is usually used when uncertainties in data exist,¹⁶ a rolling time horizon is also applicable to reduce problem size, see for example the paper by Merce and Fontan¹⁷ where a rolling horizon is used to solve capacitated lot sizing problems, and Eveborn and Rönnqvist¹⁸ where they apply a rolling horizon method on

a staff scheduling application in a starting phase, followed by an improvement phase where some time periods are selected and reoptimized. Dimitriadis et al.¹⁹ treat both a forward and backward rolling horizon method for scheduling of multipurpose plants. In Axsäter and Schneewiess²⁰ a rolling horizon approach is used to optimize a hierarchical planning system.

We have developed a mixed integer optimization model in which the scheduling decisions are modeled with binary variables for each vessel, for each possible departure time and each route. The planning covers a time horizon that is usually around one month and is operative with time periods of about one day. We use a rolling horizon and solve the model repeatedly for different short ranges, for example two weeks. Overlapping time horizons are also modelled as well as an extended range of periods with relaxed binary variables, but we do not consider the whole horizon in every solution range. The main contribution of the paper is an robust approach (both model and method) to solve large scale ship routing problem in practice. This includes to consider the many products as well as the multiple pick up and delivery aspects. Furthermore, the approach is insensitive to problem size as there is a possibility to control the size of the underlying optimization problem.

In the computational experiments we find that even short ranges (ten days) can be used to find acceptable solutions. The solution times are within 30 minutes for a one-month schedule, compared to the manual planning where more than one day is used to set up a feasible set of schedules. For several test instances the algorithm performed better than the state-of-the-art code CPLEX - taking time efficiency into account.

The outline of this paper is as follows. We give a detailed description of the problem in the next section, where we also present Södra Cell's complete supply chain. In the following two sections we give the mixed integer programming model and describe in detail the rolling horizon approach. Next follow the computational results and finally we end with some concluding remarks.

Problem formulation

Components of the supply chain

The supply chain begins in Scandinavian forests, typically in harvesting areas owned by the members of Södra. These areas are aggregated by Södra into domestic supply forest districts. The supply of raw material includes soft and hard wood from these forests, together with imports of raw materials from other countries, mainly Russia and the Baltic States. It is customary for Södra to choose domestic wood over non-member deliveries. It is possible to store harvested pulp wood within a forest district before it is transported to the production mills. The byproduct from saw mills in the form of wood chips is another important raw material used in pulp production. The domestic wood is transported to the pulp mills by truck, while imported wood is delivered with vessels.

There are five pulp mills in management by Södra, where today four mills share a common production and distribution planning. (The fifth mill produces a chemo thermo mechanical pulp while the other four produce chemical bleached pulp.) In total there are around 30 different

products and at each pulp mill only one product (pulp type) can be produced at a time (with one exception where two production lines are available). Some products can be produced in more than one mill, while other products are specifically produced at one mill.

To distribute pulp, Södra Cell uses three time chartered vessels (TC-vessels) that are hired on long term contracts. In addition, they lease spot vessels for short routes and deliver pulp directly by train or truck. Close to each mill there is a shipment port, either rented or owned by Södra. It is possible to store pulp (typically pulp destined for export) at these ports. In addition to the export customers there are domestic customers whose demands are met using deliveries by truck and train.

Although the majority of export customers are located in Western Europe, deliveries are also made to overseas customers. The wide spread of customer locations makes the distribution planning a difficult but important task. It is not only difficult because of the associated travel distances, it is also hard to meet all the demands for different products with the limited range of production possibilities available.

The overall picture of the supply chain is shown in Figure 1. In this paper we consider only the problem concerning distribution to customers, that is from the pulp stocks in Figure 1.

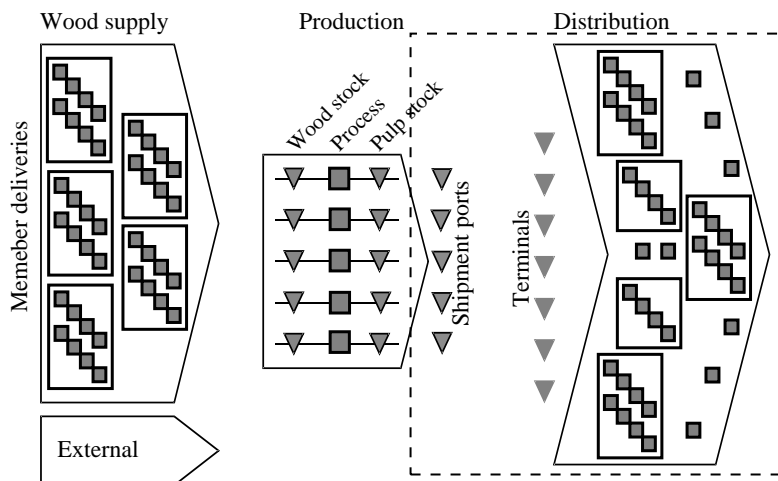


Figure 1: Södra Cell’s supply chain. The member areas (squares to the left) are aggregated to larger districts and likewise the customers (squares to the right) in delivery points, destinations.

Distribution

There are large quantities of pulp to deliver to customers located all around the world. To accomplish this, pulp is distributed by vessels, trucks and railway. In Europe, Södra Cell uses the time chartered vessels, only for this purpose. On the European market, the chain from pulp mill to customer includes: deliveries from mill to shipment port, loading at port and unloading at terminal suitable for further delivery to customers.

There are around 30 ports with terminals in Europe that are used by Södra Cell. The TC-vessels are scheduled on a continual basis where they in a schedule repeatedly visit one to three shipment ports (close to pulp mills) followed by one to three harbour terminals. The final deliveries are by barges, railway and trucks. Barges are only used to deliver products from harbour terminals to inland terminals, while trucks and railway are used all over Europe for the deliveries from terminals to the customer.

Planning Decisions

For a distribution planner there are both strategic and operative decisions. For example, to what extent should spot vessels be used? should some vessels use fixed routes (repeat the same routes in a sequence)? which terminal should be used? are examples of strategic decisions. On operative level we have a wide range of decisions, summarized in this section.

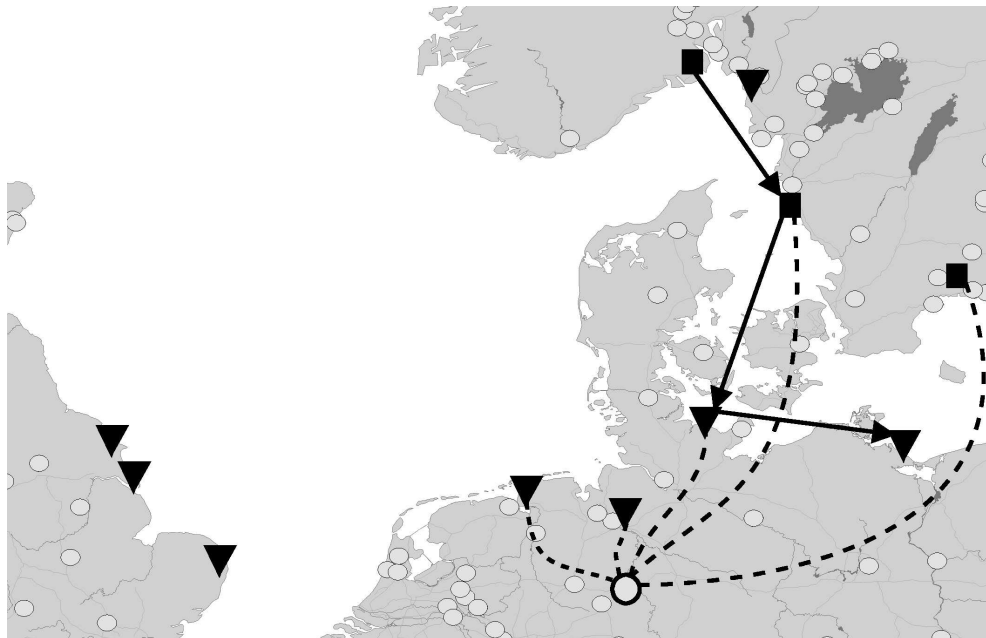


Figure 2: Different pulp flows. The squares illustrate pulp mills, the triangles harbour terminals and the circles customers. The dashed lines are flows to customer, either directly from a pulp mill or from a harbour or inland terminal, and the arrows form one route for one vessel.

In Figure 2 we see an example of different flows to one customer. Due to the variety of delivery options and the options' dependency on product and destination, the routing and scheduling decisions cannot easily be extracted from the supply chain context and be planned separately. A route in this context is a sequence of shipment ports (loading only) followed by a sequence of harbour terminals in ports for unload purpose only. The scheduling process is to assign routes for the TC-vessels and to rent extra vessels for single routes over a time horizon. A summary of decision,

- Routing and scheduling of the TC-vessels
- When and where a spot vessel is used
- Size of spot vessel
- Direct flows from pulp mill to customer
- Choice of supply mill for product
- Storage issues in terminals
- Usage of inland terminals
- Quantities of pulp products on vessels

Normally a four-five weeks long schedule for the TC-vessels is maintained and revised every week. When the TC-vessels are scheduled, the need for the additional spot vessels is identified. If the sales term is such that Södra Cell is responsible for arranging the forwarding to the customer, then the final delivery planning is done following the ship schedules. In the next section we formulate an mathematical model which considers these steps simultaneously.

MIP model

To begin with, we define a set V as the enumeration of all vessels where V_S is the subset of spot vessels and V_C is the subset of TC-vessels. Other sets used in the formulation are, P for products, C for customer locations (destinations), T for time periods, M for shipment ports (ports close to pulp mills), H for harbour terminals (terminals in a port for unloading purpose only), L for inland terminals (not located in a port) and we define $J = H \cup L$.

Define a route to be a sequence of shipment ports followed by a sequence of harbour terminals. All the routes are pre-generated and enumerated in a set R . A stop in a shipment port is considered as a trivial route, with length one time period (typically one day). A vessel can be on stand by, but only in a shipment port.

A number of modeling sets and parameters are used to formulate the scheduling constraints and the upper bounds of route flows. We assign an index for every combination of shipment port and terminal within a route. These indices are unique for each pair of route and time period, that is for each start period for a route throughout the time horizon. We define the set I_{jt} , with $j \in M \cup H$ and $t \in T$, as the set of all indices that correspond to the feasible flow links in time period t that a vessel can have from a shipment port, or to a terminal j via a route. In Figure 3 we see an illustration of the index structure.

In the figure, the index i on the arc from (t, m_1) to $(t+3, h_2)$ is included in set $I_{m_1 t}$ and the set $I_{h_2, t+3}$. These indices together determine all possible vessel flows. To associate the indices with the routes, we use a set I_{rt}^R , the set of indices that corresponds to flow on route r if the route begins




Figure 3: Route indices. The vertical axis shows ports, m_1 to m_3 are shipment ports and h_1 to h_3 are terminals. The horizontal axis has time periods. r_1 and r_2 are two routes and the dashed lines show possible flows.

in time period t . In Figure 3, the indices for the upper two dashed lines would be the members of the set $I_{r_1, t-1}^R$ and the indices for the lower four dashed lines the members of the set $I_{r_2, t-1}^R$.

Define $L(r)$ to be the length of the route r measured in time periods and $D(r', r)$ to be the number of time periods it takes for a vessel from the last terminal of route r' to the first shipment port of route r , assuming that all vessels have the same velocity. Then to model the number of time periods between two routes we use a help set, $T_{trr'}$ for the interval $\max\{p_0, t - L(r') - D(r', r) + 1\}, \dots, t - 1$, where p_0 is the first time period in the whole time horizon. The set $T_{trr'}$ is used to forbid two routes to be used simultaneously by one vessel. In other words, $T_{trr'}$ is the set of time periods that a vessel will need from the beginning shipment port of route r' to the beginning shipment port of route r , if the route r starts in time period t . See Figure 4 for an example.

Note that we in both the definition of I_{jt} and the definition of $T_{trr'}$ assume that all vessels have the same velocity. This is for simplicity in notation, differences in velocity can be taken into account if we instead define these sets (and the functions L and D) with respect to vessel or vessel properties.

We use continuous variables to model storage volumes in terminals, shipment ports and customer locations, direct flows from pulp mills to customers, flows between harbour terminals and inland terminals and flows from terminals to customers. These variables are defined as




Figure 4: Illustration of the set $T_{trr'}$. In the figure two different routes are shown. Each bar is a route starting at a certain time. If a vessel begins route r in time period t , then it cannot use the route r' in any period $t' \in T_{trr'} = \{t-3, t-2, t-1\}$.

- $l_{jpt}^H, l_{jpt}^L,$
 l_{jpt}^M, l_{jpt}^C = quantity stored in location j of product $p \in P$ at (the end of) time period $t \in T$ in harbour terminal (H), inland terminal (L), shipment port (M) or in customer location (C)
- y_{jrpt} = direct flow from shipment port $j \in M$ to customer $r \in C$ with product $p \in P$ in time period $t \in T$
- u_{krpt} = flow from terminal $k \in J$ to customer $r \in C$ of product $p \in P$ in time period $t \in T$
- w_{hkpt} = flow from harbour terminal $h \in H$ to inland terminal $k \in L$ of product $p \in P$ in time period $t \in T$

The variables $l_{jp0}^H, l_{jp0}^L, l_{jp0}^M$ and l_{jp0}^C are the initial storage volumes respectively, in location j of product $p \in P$. There is one binary variable in the model for every vessel (including the spot vessels), route and time period in combination. Let I denote the set of all route flow indices and define the route variables as

- z_{srt} = 1 if vessel $s \in V$ uses route $r \in R$ in time period $t \in T$, 0 otherwise
- x_{sip} = flow on vessel $s \in V$, index $i \in I$ with product $p \in P$

There are parameters for supply, demand, capacities:

- S_{jpt} = supply of product $p \in P$ in shipment port $j \in M$ in time period $t \in T$
 D_{rpt} = demand of product $p \in P$ by customer $r \in C$ in time period $t \in T$
 B_k = handling capacity in shipment port / terminal $k \in M \cup J$ per time period
 K_s = shipping capacity of vessel $s \in V$

The constraints are as follows,

$$\begin{aligned}
l_{jpt}^M &= l_{jp,t-1}^M - \sum_{s \in V} \sum_{i \in I_{jt}} x_{sip} \\
&\quad - \sum_{r \in C} y_{jrpt} + S_{jpt}, \quad j \in M, p \in P, t \in T
\end{aligned} \tag{1}$$

$$\begin{aligned}
l_{jpt}^H &= l_{jp,t-1}^H + \sum_{s \in V} \sum_{i \in I_{jt}} x_{sip} \\
&\quad - \sum_{r \in C} u_{jrpt} - \sum_{k \in L} w_{jkpt}, \quad j \in H, p \in P, t \in T
\end{aligned} \tag{2}$$

$$\sum_{s \in V} \sum_{i \in I_{kt}} \sum_{p \in P} x_{sip} \leq B_k, \quad k \in H \cup M, t \in T \tag{3}$$

$$\sum_{h \in H} \sum_{p \in P} w_{hkpt} \leq B_k, \quad k \in L, t \in T \tag{4}$$

$$l_{kpt}^L = l_{kp,t-1}^L + \sum_{h \in H} w_{hkpt} - \sum_{r \in C} u_{krpt}, \quad p \in P, k \in L, t \in T \tag{5}$$

$$l_{rpt}^C = l_{rp,t-1}^C + \sum_{j \in M} y_{jrpt} + \sum_{j \in J} u_{jrpt} - D_{rpt}, \quad r \in C, p \in P, t \in T \tag{6}$$

$$\sum_{i \in I_{rt}^R} \sum_{p \in P} x_{sip} \leq K_s z_{srt}, \quad s \in V, r \in R, t \in T \tag{7}$$

$$\sum_{r \in R} z_{srt} \leq 1, \quad s \in V_C, t \in T \tag{8}$$

$$z_{srt} = 1 - \sum_{r' \in R} \sum_{t' \in T_{rr'}} z_{sr't'}, \quad r \in R, s \in V_C, t \in T \tag{9}$$

We have flow conservation in the shipment ports (1) and flow conservation in the harbour terminals (2). Next follows capacity constraints for harbour terminals and shipment ports (3) and capacity constraints at inland terminals (4). There are flow conservation constraints for inland terminals in (5) and flow conservation in customer destinations, (6). The variable u_{khpt} in (5) and the variables y_{jrpt} and u_{jrpt} in (6) can in the constraints be delayed with the delivery time, although it is (for simplicity in notation) left out in the formulation above. The route scheduling definitions are comprised in the constraints (7), (8) and (9). In (7) the constraints simply allow flow up to maximum capacity for vessel s on route r in time period t , if the route is used, that is if $z_{srt} = 1$. All vessels, including spot vessels comply with the constraints in (7). The constraints in (8) treat the case that a vessel can not begin more than one route in the same time period. These constraints together with the constraints in (9) force the vessels in V_C to use a route in every time period, either on water or ported in a shipment port on a trivial route. The vessel s in (9) either begins the route r in period t or began a route $r' \in R$ during the period in $T_{trr'}$.

The objective is to minimize the sum of all costs. For TC-vessels, there is a fixed daily cost, a distance cost and a loading or unloading cost associated with each port. Since the route ends in a terminal, a cost for the return to a shipment port is treated separately. The total return cost is expressed as

$$\sum_{s \in V_C} c^{return} \left(|T| - \sum_{r \in R} \sum_{t \in T} L(r) z_{srt} \right)$$

where c^{return} is the cost per period on water and $|T|$ is the length of the whole time horizon. The total number of return periods is expressed within the bracket. The problem is

$$\begin{aligned} \min \quad & \sum_{j \in M} \sum_{p \in P} \sum_{t \in T} c_j^M l_{jpt}^M + \sum_{j \in H} \sum_{p \in P} \sum_{t \in T} c_j^H l_{jpt}^H + \sum_{j \in L} \sum_{p \in P} \sum_{t \in T} c_j^L l_{jpt}^L \\ & + \sum_{j \in C} \sum_{p \in P} \sum_{t \in T} c_j^C l_{jpt}^C + \sum_{j \in M} \sum_{r \in C} \sum_{p \in P} \sum_{t \in T} c_{jr}^Y y_{jcpt} + \sum_{k \in J} \sum_{r \in C} \sum_{p \in P} \sum_{t \in T} c_{kr}^U u_{kcpt} \\ & + \sum_{h \in H} \sum_{k \in L} \sum_{p \in P} \sum_{t \in T} c_{hk}^W w_{hkpt} + \sum_{s \in V} \sum_{i \in I} \sum_{p \in P} c_{si}^X x_{sip} + \sum_{s \in V} \sum_{r \in R} \sum_{t \in T} c_{sr}^Z z_{srt} \end{aligned}$$

s.t. the constraints (1) – (9)

all variables ≥ 0

$z \in \{0, 1\}$

The coefficients c_j^M , c_j^H , c_j^L and c_j^C are dependent on if the storage facilities are owned by Södra Cell, where they are located and approximations of changes in market prices. The c_j^Y , c_j^U and c_{hk}^W are estimated from historical data and depend on distances and delivery agreements. The fixed costs for routes, c_{sr}^Z , are determined by port contracts, distances (with the return costs) and vessel dependent costs. The coefficient c_{si}^X includes the loading and unloading cost for the pulp volume and is dependent on port and vessel contracts.

Solution Method

The idea of the algorithm is to repeatedly solve a limited MIP using a branch and bound search algorithm. Each limited MIP comprises of a short time horizon, and the overall problem is solved when all time periods have been considered in at least one limited MIP.

In practise we begin with a selection of an initial set of periods. To this solution range we add an extension, where the binary variables are relaxed, that serves as a forecast range. Ideally we would like to include the full remaining period in order to get information about the demand over the entire planning period. In most cases, however, the demand profile is rather stable and we can keep it down and focus on avoiding problems (with low future flexibility at the integer solution to be fixed) in a local sense. Over the current horizon (the solution range with the forecast range), we solve the underlying MIP using branch and bound. Then an in advance defined part of the obtained solution, the frozen periods, is saved and considered as determined. We continue by moving the solution range to a new set of time periods. The new solution range extends the horizon with new time periods but keeps a common interval with the previous range (all but its frozen periods). In Figure 5 we illustrate a schedule for one vessel and the overlapping horizons.

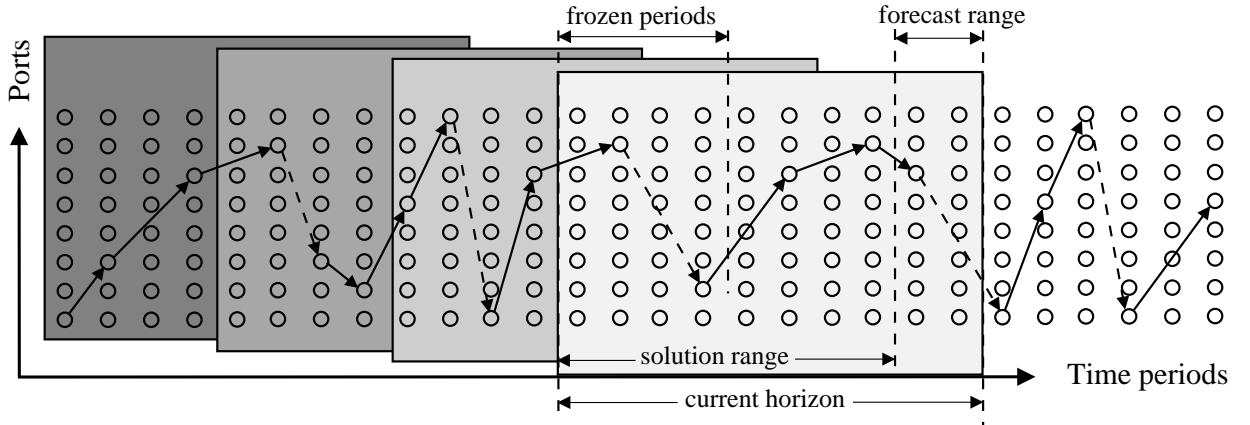


Figure 5: The rolling time horizon for an example of a schedule for one vessel.

The frozen periods define new initial storage volumes. The routes that begin in the overlapping periods are set free, non fixed, while the routes being fixed define a starting position for the vessels, see Figure 6. In the figure the route r_1 is not fixed, while r_2 and r_3 are considered as determined. In the case of r_1 it is necessary to model frozen periods from the previous range in the current horizon. In the model this is included by letting the constraints in (9) consider all the relevant time periods.

The repeated solving continues until the whole planning horizon is covered. Finally, keeping the (binary) route variables fixed, we solve the linear program for the whole horizon to find optimal flows.

All routes are generated in advance by looping through all combinations of shipment ports and harbour terminals (with a maximum number of ports and terminals). Depending on the data set,




Figure 6: Routes that overlap between the solution ranges.

for instance what the demanded products from one particular customer are, the model size can be significantly reduced. Therefore a preprocessing phase initializes the model. It includes reduction of constraints and variables without model simplifications. The reduction handles questions like,

- Only available products at a pulp mill throughout the horizon can be supplied from the mill.
- Only products demanded from a customer can be transported to the customer.
- A terminal can only handle products that can be delivered to a customer from the terminal.
- A route can only carry products that can be supplied from one of the visited mills and delivered to one of the visited harbour terminals.

Unexpected infeasibilities and inconsistencies in data are treated with dummy variables d_{rpt} , dummy flow to customer r with product p in time period t . The dummy variables are included with high cost in the objective function and in the flow conservation constraints for customers, (6). There are two main reasons to use these variables. First, and most important, is that, in practice, there often, if not always, are errors in the large set of data coming from the company database. By having these dummy variables makes it very easy to identify these problems so that they can be fixed. Second, in our tests we sometimes got a solution that was slightly infeasible. In all cases the reason was that a very small volume of some product was missing. In practice this is easily fixed once the solution process terminates as the supply and demand only are estimates. Also, in practice trucks can be used for quick deliveries.

Algorithm summary

Step 0 Set a period length and generate the set of routes. Initialize the algorithm with the last time period of the whole time horizon P_N , a first start period p_0 and the solution range T with length p . Let $p_1 := p_0$ be the first period of the solution range. Define F as the set of frozen periods and let q be the number of periods to fix in each iteration.

- Step 1** Solve the MIP over $T = \{p_1, \dots, p_1 + p - 1\}$ with the additional forecast range
- Step 2** Fix variables in $F = \{p_1, \dots, p_1 + q - 1\}$
- Step 3** If $p_1 + p + q > P_N$ then goto step 4, else let $p_1 = p_1 + q$, update T and F and goto step 1
- Step 4** If $p_1 + p - 1 = P_N$ then goto step 5, else let $p_1 = P_N - p + 1$, update T and F and goto step 1
- Step 5** Unfix all flow variables and solve LP with $T = \{p_0, \dots, P_N\}$. Stop.

The most time consuming step is step 1, where the MIP is solved, either to optimality or until a stop criteria (MIP-gap or time limit) is fulfilled. In Step 5 all routes are fixed and only an LP problem is solved.

The condition in step 3 is true if the current or the next solution range is the final range. In step 4, if the condition $p_1 + p - 1 = P_N$ is not satisfied, we include the remaining time periods. Here the forecast range can either be included, if data is available, or be excluded completely. There is no guarantee for optimal nor feasible solutions as soon as the solution range is shorter than the whole time horizon.

Computational Results

Five problem instances have been used for computational experiments and a total number of 60 runnings were performed using different parameters for the instances. We followed Södra's manual planning horizon and set the time discretization in all instances to daily level (24 hours) and the whole horizon length to 40 days. The algorithm is implemented in AMPL version 10.6 with the solver CPLEX 7.0 and we used a Pentium 4, 1700 MHz, 1 GB RAM computer running windows NT. In our implementation we set up the problem for the whole time horizon in AMPL, though we considered only the current horizon while solving. To include the whole horizon had the advantage to let us calculate the objective function value of the LP relaxation.

Cases

Most of the data was taken from a real world case (computed by staff at Södra Cell) from the planning period April to June 2003. Some of the spot vessel costs were estimations, including costs for flows on spot routes. Initial storages were estimated from an average storage volume and all vessel velocities were set to the average speed of the TC-vessels. The demand was approximated from the actual deliveries made by Södra during the time considered.

Depending on the time discretization different capacities apply to the ports. To load a vessel with a maximum capacity of 5600 tonnes takes around eight hours. With two feasible eight hour shifts in a port, and assuming that only one vessel can be ported at the time, this gives an upper bound of 11 200 tonnes of pulp per day. In all instances we had two spot vessels defined, with a

maximum capacity of 2600 tonnes of pulp. They were modeled as TC-vessels without return costs to restrict the number of available spot vessels.

All routes were generated as a list of ports and every port was given a time offset relative to the first port in the route. The length of a route (in average 3 to 10 days) was set to the sum of travel time, loading and unloading time at ports. This estimated time was rounded up to nearest number of time periods. Due to this, a route never visited two ports in the same time period.

The instances are shown in Table 1 as well as the size of the problems in terms of variables and constraints. The column “variables” includes the binary variables. The numbers in all columns refer to the problem modeled for the whole planning horizon.

Instance	$ R $	$ J $	$ S $	$ C $	$ P $	$ M $	#Binaries	#Variables	#Rows
1a,1b,1c	27	3	2	101	12	3	3 984	36 890	10 500
2a,2b,2c	36	3	2	101	12	3	5 240	58 318	12 770
3	360	8	2	110	12	4	62 630	478 557	128 341
4	660	11	2	118	12	4	112 090	873 397	227 798
5	1253	31	3	269	30	4	211 340	1 273 194	431 611

Table 1: Problem instances. The columns are; number of routes ($|R|$), terminals ($|J|$), TC-vessels ($|S|$), customers locations ($|C|$), products ($|P|$) and mills ($|M|$). The number of variables and rows for the instances 1a, 1b and 1c is the average number, as well as for 2a, 2b and 2c.

The instances 1 to 4 are based on a reduced set of data, with no differences between the instances 1 and 2 except in the number of routes modeled. For the instances 1a-c we generated direct routes (one shipment port to one terminal) and all combinations of two shipment ports and one terminal. The instances 2a-c also include the combinations of one shipment port and two terminals. The third and the fourth instances are slightly larger. They have more ports, some routes with four ports and more customers. The fifth instance correspond to the real world case and includes the routes with up to four ports in sequence. Among the routes with two terminals we included only the combination that gave the shortest distance, measured in time periods. To take the limited supply into account we let all combinations of shipment ports be included and in the case that more than one combination gave the same time, both were included.

In the problem instances one to four, there are two TC-vessels scheduled on long term contract whereas in the instance five all three TC-vessels are included.

Solution quality

The instances 1 and 2 have been solved with three different sets of data, selected from different planning periods. In Table 2 these planning periods are numbered a to c . The table displays the objective function values obtained with different parameters for the algorithm. Row LBD is the best known lower bound for the problem for the whole horizon. The row (40/*/1) is the solution from the branch and bound search over the whole time horizon. This correspond to solve the whole model in one instance. The following rows are the objective function values obtained for different parameter settings (Le/Fr/Gap), in order, the length of the solution range (excluding the five forecast periods), the number of frozen periods and the maximum MIP-gap accepted. Some

parameter settings gave infeasible solution marked with *inf* in the table. The instance 5 was only solved with the maximum MIP-gap 4%. In this case we could not solve the entire model directly.

(Le/Fr/Gap)	1a	1b	1c	2a	2b	2c	3	4	5
40/*/1	2 216	2 213	2 332	2 199	2 208	2 330	3 543	4 767	-
20/5/4	2 263	2 304	2 428	2 319	2 380	2 469	3 672	4 881	13 367
20/5/1	2 246	2 247	2 335	2 250	2 290	2 331	3 635	4 897	-
20/10/4	2 237	2 253	2 365	2 357	2 319	2 372	<i>inf</i>	4 812	13 254
20/10/1	2 223	2 198	2 279	2 253	2 254	2 296	<i>inf</i>	4 815	-
10/5/4	2 216	2 252	2 298	2 212	2 273	2 348	3 506	<i>inf</i>	13 218
10/5/1	2 214	2 229	2 285	2 223	2 199	2 287	3 501	4 805	-
LBD	2 181	2 179	2 256	2 178	2 177	2 253	3 489	4 719	13 111

Table 2: Objective function values for different problem instances and parameter settings. The first row (40/*/1) is the objective from the the branch and bound search on the problem instances over the whole time horizon. The last rows (LBD) is the objectives of the strengthened LP relaxation. The *inf* means that dummy variables were active in the solution found within time limit.

Better results were obtained with the short solution range of 10 time periods and a 4% MIP-gap than with the 20 periods range with a 4% MIP-gap for all instances, except for the fourth instance where no feasible solution was obtained with the setting. In the cases when no solution were found it may be due to the fact that a rather large proportion of the horizon was frozen together with relatively long routes. Compared to the settings with a 1% gap, the solution quality is higher with a lower gap, and for the instances 1b and 1c the setting 20/10/1 obtained the overall best solution. For only two instances the B&B search over the entire horizon obtained the best solution. Cases 2a, 2b, 2c have larger solution space (more routes) as compared to 1a, 1b and 1c respectively. The fact that the solutions from case 1 often have better solutions comes with the fact that the solver stops directly once the convergence tolerance in the solver is met. The same argument holds for the fact that row (40/*/1) theoretically should have the best solution.

Solution time

The time limit for each branch and bound search was set to one hour, except for the whole horizon, which was limited to two hours (or a MIP-gap of one percent). Table 3 shows the solution time in seconds used by the algorithm for every instance and parameter setting. The italic entry in every column is the time for the setting that obtained the best solution found. The settings where one or more branch and bound search exceeded the time limit is shown in the table as >2 hours.

	1a	1b	1c	2a	2b	2c	3	4	5
40/*/1	2h	2h	2h	2h	2h	2h	2h	<i>5 130</i>	
20/5/4	140	45	18	385	188	30	815	1 432	1 655
20/5/1	739	657	102	1 562	762	213	1 338	1 931	-
20/10/4	108	49	37	57	176	46	>2h	1 551	2 383
20/10/1	384	<i>4 120</i>	<i>4 926</i>	632	>2h	>2h	>2h	3 447	-
10/5/4	19	17	25	23	28	50	2 710	1 044	<i>1 485</i>
10/5/1	<i>146</i>	520	1 272	267	<i>1 314</i>	<i>1 996</i>	>2h	<i>1 950</i>	-

Table 3: Solution time in seconds for each instance and parameter group.

There were large differences between the time requirements for the instances $1a$, $1b$ and $1c$ and between $2a$, $2b$ and $2c$. That is, the variation in solution time was dependent on the problem size and also on the considered planning period. The setting 20/10 compared to 20/5 was more time consuming, even though the number of evaluated horizons was less.

Concluding remarks

We have proposed a mixed integer programming model for a difficult combined distribution and ship scheduling problem. The problem includes many products and multiple pick up and deliveries. The model is chosen such that we can identify errors in the data and always guarantee a near feasible solution. Based on feedback from the manual planners we know that the model accurately describes the actual real problem.

The solution method is based on a rolling time horizon and branch and bound to find good feasible schedules and optimal pulp flows for large scale practical problems. The approach is insensitive to problem size as there is a possibility to control the size of the underlying optimization problem. The computational experiments are worked out on a set of small test cases that are compared with a direct approach with branch and bound, and on a full scale real world case.

The formulated problem with the suggested solution approach was solved and we obtained acceptable solutions within reasonable time limits. The algorithm is very efficient compared to manual planning. Although there was no guarantee that the suggested algorithm finds solutions, the computational results indicated that even with very short solution ranges (in the algorithm) feasible schedules are obtained.

The selection of parameters in the approach such as length of a time period, current horizon, frozen period, forecast range is depending on a number of aspects. Some are route lengths, variations in demand and supply and type of operations. It is therefore difficult to establish a general set of parameter values. Instead these need to be adjusted for each application. This is however not a difficult issue as it is rather easy to find a suitable set of values (and there is only a few parameters) once the application and case is known.

An interesting future study involves the possibility to speed up the solution process by starting with a limited set of potential routes and once a feasible one is found the set is dynamically increased.

References

- 1 J.F. Shapiro (ed) (2001) *Modeling the Supply Chain*. Duxbury
- 2 A. Drexler and A. Kimms (1997). Lot sizing and scheduling - survey and extensions, *European Journal of Operational Research* **99**: 221–235
- 3 M. Christiansen, K. Fagerholt and D. Ronen (2004). Ship Routing and Scheduling: Status and Perspectives, *Transportation Science* **38**: 1–18

- 4 L.H. Appelgren (1969). A column generation algorithm for a ship scheduling problem, *Transportation Science* **3**: 53–68
- 5 L.H. Appelgren (1971). Integer programming methods for a vessel scheduling problem, *Transportation Science* **5**: 64–78
- 6 K. Fagerholt (2001). Ship scheduling with soft time windows - An optimisation based approach, *European Journal of Operational Research* **131**: 559–571
- 7 K. Fagerholt and M. Christiansen (2000). A combined ship scheduling and allocation problem, *Journal of the Operational Research Society* **51**: 834–842
- 8 D. Carlsson and M. Rönnqvist (1999) Wood flow problems in Swedish forestry, Edited by G. Frumerie, Report No. 1, The Forestry Research Institute of Sweden, 48 pages, Sweden
- 9 D. Bredström, J.T. Lundgren, R. Rönnqvist, D. Carlsson and A. Mason (2004). Supply chain optimization in the pulp mill industry - IP models, column generation and novel constraint branches, *European Journal of Operational Research* **156**: 2–22
- 10 D. Bredström, J.T. Lundgren, M. Rönnqvist, D. Carlsson and A. Mason (2001). Supply chain optimization in the pulp mill industry, Proceedings of the Logistics Research Network 6th Annual Conference, Heriot-Watt University, 13-14 September, 67–74
- 11 W.J. Bell, L.M. Dalberto, M.L. Fisher, A.J. Greefield, R. Jaikumar, P. Keida, R.J. Mack and P.J. Prutzman (1983). Improving the Distribution of Industrial Gases with an On-line Computerized Routing and Scheduling Optimizer, *Interfaces* **13**: 4–23
- 12 R.C. Larson (1988). Transporting Sludge to the 106-Mile Site: An Inventory/Routing Model for Fleet Sizing and Logistics System Design, *Transportation Science* **22** 186–198
- 13 E.J. Beltrami and L.D. Bodin (1974). Networks and Vehicle Routing for Municipal Waste Collection, *Networks* **4**: 65–94
- 14 M. Christiansen (1999). Decomposition of a Combined Inventory and Time Constrained Ship Routing Problem, *Transportation Science* **33** 3–16
- 15 D.E. Blumenfeld, L.D. Burns, C.F. Daganzo, M.C. Frick and R.W. Hall (1987). Reducing Logistics Costs at General Motors, *Interfaces* **17**: 26–47
- 16 B. Karami, S.M.T. Fatemi Ghomi and J.M. Wilson (2003). The capacitated lot sizing problem: a review of models and algorithms, *Omega* **31**: 365–378
- 17 C. Merce and G. Fontan (2003). MIP-based heuristics for capacitated lotsizing problems, *International Journal of Production Economics* **85** 97–111
- 18 P. Ekebom and M. Rönnqvist (2004). Scheduler – A system for staff planning, *Annals of Operations Research*, **128**, 21–45
- 19 A.D. Dimitriadis, N. Shah and C.C. Pantelides (1997). RTN-based Rolling Horizon Algorithms for Medium Term Scheduling of Multipurpose Plants, *Computers chem. Engineering* **21**: 1061–1066
- 20 S. Axsäter and C. Schneewiess (eds) (1986) Hierarchical production planning: Tuning aggregate planning with sequencing and scheduling, Multistage production planning and inventory control. Berlin