

KATHOLIEKE
UNIVERSITEIT
LEUVEN

DEPARTEMENT TOEGEPASTE ECONOMISCHE WETENSCHAPPEN

RESEARCH REPORT 0005

A TIRE PRODUCTION SCHEDULING SYSTEM FOR SOLIDEAL

by
Z. DEGRAEVE
R. JANS

D/2000/2376/05

KATHOLIEKE UNIVERSITEIT LEUVEN

Department of Applied Economics

A TIRE PRODUCTION SCHEDULING SYSTEM FOR SOLIDEAL

Zeger Degraeve

London Business School, London, U.K.

Katholieke Universiteit Leuven, Department of Applied Economics, Leuven, Belgium

Raf Jans

Research Assistant of the Fund for Scientific Research - Flanders (F.W.O.)

Katholieke Universiteit Leuven, Department of Applied Economics, Leuven, Belgium

Zeger Degraeve
London Business School
Sussex Place, Regent's Park
London NW1 4 SA, U.K.
email: zdegraeve@london.edu

Raf Jans
Research Assisant of the Fund for Scientific Research – Flanders (F.W.O.)
Katholieke Universiteit Leuven
Department of Applied Economics
Naamsestraat 69, 3000 Leuven, Belgium
email: raf.jans@econ.kuleuven.ac.be

January 17, 2000

A TIRE PRODUCTION SCHEDULING SYSTEM FOR SOLIDEAL

Zeger Degraeve
Raf Jans

Abstract

Solideal is an international industrial tire manufacturer. The tires are built in molds and are produced in heaters in large series production runs. Some tires can be cured in two different types of heaters. Preparing a heater for a specific tire type requires a set up which typically only takes place during the morning shift. Therefore, the company employs to a discrete production policy. The resulting scheduling problem is a capacitated, multi-item lot sizing problem with set up times and multiple non-identical alternative machines. The objective is the minimization of the total costs for set up, backlog and inventory. Management opted for a fast algorithm that could produce various alternative high quality solutions. The problem is solved using a column generation scheme where the master is solved approximately using Lagrange relaxation. During the column generation process, various heuristic integer solutions are created from which the schedulers can then choose one to implement. We present computational results on real life data with up to 30 products and 30 periods.

Keywords: *Lot Sizing; Production Planning; Column Generation; Lagrange Relaxation; Decomposition; Set Up Times.*

Introduction

Solideal is an international industrial manufacturer with a yearly sales figure of $150 million. The group manufactures and distributes a whole range of products: solid industrial tires, pneumatic industrial tires, steel products and rubber tracks. These products are used on a variety of industrial and construction equipment. Their applications vary from factories to airports, harbors, building sites and quarries. The tires make up more than ninety percent of the total sales volume. They are mainly produced in three factories in Colombo, Sri Lanka and, from there, they are distributed all over the world. The choice for Sri Lanka was predicated upon the following factors: (1) the previous experience of one of the founders establishing a similar venture there; (2) the availability of skilled manpower at low cost; (3) the availability of locally produced rubber and (4) government investment incentives. Currently Solideal has more than 4000 employees and produces about 475 different tires. In fifteen years the Solideal group has managed to become one of the major manufacturers and distributors of industrial tires worldwide. This was achieved by their ability to manufacture a good quality product which could be sold at a competitive price. Solideal is focussing on niche markets in the tire business because this allows larger margins while not being attacked by the big tire manufacturers who find these markets too small. The current strategy is still aimed at expansion. Production should double within the next three to five years. In order to achieve this growth, substantial investments in new capacity are required. Management believes that one of the cornerstones for the success of the growth strategy will be the

efficiency of the production planning and logistics system. Solideal was facing both substantial backlog and inventories. The management came to us with the challenge to make a scheduling system which (1) could produce high quality production plans, (2) was fast enough to be used several times consecutively on a daily basis, (3) where the planners had still control over the important parameters, so that by varying them they would be able to generate alternative schedules and (4) where the limitations of the production environment were appropriately taken into account.

In Section 1 the problem is described. In Section 2 we give a review of the related literature. Section 3 describes our algorithm. The problem is solved using a column generation scheme, where the subproblems are optimized with a dynamic programming (DP) recursion. The master is solved either approximately using Lagrange relaxation and subgradient optimization or exactly using LINDO. We further discuss the calculation of upper bounds and lower bounds. Next we discuss the issue of fast movers and other industrial complications in Section 4. Section 5 gives an overview of the computational results. Finally Section 6 presents the model impact.

1. Problem description

Solideal faces a complex production planning process. Several hundreds of different tires must be scheduled. The production is based on a JIT manufacturing process. A small stock of green tires is available. A green tire consists of several layers of fabric and rubber and is still flexible. At the final stage of the production process, the green tires are put in molds and the molds are placed in heaters to cure the tire. The curing process is pulled by the demand. Each heater contains exactly one mold at the time. Solideal disposes of several types of regular heaters and several types of BOM heaters. The BOM heaters can withstand higher pressure and temperatures, resulting in a shorter curing time. A mold can only be used for one specific tire-heater combination. For each different tire-heater combination a different kind of mold is needed. The cycle time for a specific tire-heater combination is the total of the curing time and handling time. The handling time is the time necessary to take a mold out of the heater, to remove the tire, to put a new green tire in the mold, to press the mold and to put it back in the heater. A tire can be cured in a specific type of regular heater, in a specific type of BOM heater or in both. The fact that some tires can be produced in two different types of heaters with different efficiencies complicates the problem because all the tires sharing a same resource must be scheduled simultaneously. Setting up a heater for a new mold takes quite some time because the cold mold must be preheated. Preheating typically takes 4 to 8 hours for a regular heater and 8 hours for a BOM. Once the mold is hot, tires are produced in a continuous run. This set up time is typical for their manufacturing environment with long series production. There was a further restriction concerning the set ups. It is company practice that a set up can only take place at the beginning of the morning shift because only then qualified personnel and adequate supervision is available. This means that if we use a heater to cure a tire, we produce that tire at full capacity that day on that heater. This type of production is referred to as 'all-or-nothing' production (Schrage, 1982) or 'discrete production policy' (Magnanti and Vachani, 1990). Due to the substantial loss of time and capacity, the number of mold changes should be limited and large production runs are desired. However, the heater capacity is limited and sometimes a mold change is necessary in order to meet demand. Backlogging is allowed but must be kept at a minimum. Due to limited heater and mold capacity, every production plan will

therefore have a trade-off between the number of mold changes and backlogging. To minimize the number of mold changes, Solideal kept on producing tires for which no orders were outstanding in anticipation of future demand. This seemed to be the essential reason for the substantial backlog of tires with demand and inventories of tires without demand. They were actually producing the wrong tires. Management realized that more frequent mold changes would be required to match production with actual demand allowing to reduce both backlog and inventory. The inventory holding cost is approximately known and considered to be very small in comparison with the backlog and set up costs. The backlog cost is hard to quantify and potentially very large as the customer could become very dissatisfied not receiving his order on time. The set up cost is largely determined by the loss of capacity and as such depends on the backlog cost. By varying the relative proportions between set up and backlog cost, while keeping the holding cost very small, alternative production schedules can be generated. In this way, the trade-off between backlog and number of set ups is explored. In the end, the planners are responsible for evaluating the different production proposals.

A small example will illustrate the production process. We consider a problem with 5 products, 12 periods and 3 heater types (H1, H2 and BOM). Table 1 gives a summary of the necessary data. For each item we give the heater type(s) on which it can be produced, the number of available molds, the production rate and the set up loss in units of production. Remark that the production rate and loss per day for a specific tire-heater combination are given as an integer number. The data concerning the BOM molds are after the semi-colon. We take into account the running molds at the beginning of the horizon and assume no beginning inventory.

Table 1 Summary data for small example

Item	Heater type	Nr. molds	Prod. rate	Set up loss
1	H1; BOM	2; 3	18; 22	4; 10
2	H2; BOM	2; 3	14; 18	3; 8
3	H1	3	14	3
4	H1	2	16	4
5	H2	3	16	3

Table 2 depicts a feasible heater schedule. For every heater type there are 3 heaters available. For each period it is indicated which item is produced on what heater. Period zero gives the information on the running molds at the beginning of the horizon. Period one is the first day of the new schedule. The periods with a mold change are shaded. For example, the second heater of H2 produces tire 2 during the first 4 periods and then switches to tire 5 for 5 periods.

Table 2 Feasible heater schedule

	Period	0	1	2	3	4	5	6	7	8	9	10	11	12
H1	**1**	3	3	3	3	3	3	3	3	3	3	3	3	-
	2	3	4	4	4	4	4	4	4	4	4	4	-	-
	3	4	4	4	4	4	4	4	4	4	4	4	-	-
H2	**1**	2	2	2	2	2	2	2	2	2	2	-	-	-
	2	2	2	2	2	2	5	5	5	5	5	-	-	-
	3	-	-	5	5	5	5	5	5	5	5	5	5	-
BOM	**1**	1	1	1	1	1	1	1	1	1	1	1	1	1
	2	1	1	1	1	1	1	1	1	1	1	1	1	-
	3	1	2	2	1	1	1	-	2	2	2	2	2	2

Table 3 translates the heater schedule from Table 2 into units produced. For each tire we give the demand, production and inventory position. The periods with backlog are shaded. From Table 2 and 3 we can clearly understand the basic trade-off in our problem: in order to have fewer mold changes, we will have to allow more backlogging and vice versa.

Table 3 Demand, production and inventory position for heater schedule from Table 2

	Period	1	2	3	4	5	6	7	8	9	10	11	12
TIRE 1	Demand	0	40	30	0	170	0	0	95	40	15	80	60
	Production	24	44	56	66	66	44	44	44	44	44	44	22
	Inventory	24	28	54	120	16	60	104	53	57	86	50	12
TIRE 2	Demand	0	21	43	73	0	25	0	55	0	47	0	39
	Production	32	46	28	28	14	14	24	32	32	18	18	18
	Inventory	32	57	42	-3	11	0	24	1	33	4	22	1
TIRE 3	Demand	0	19	9	0	16	39	0	18	0	21	21	0
	Production	11	14	14	14	14	14	14	14	14	14	14	-
	Inventory	11	6	11	25	23	-2	12	8	22	15	8	8
TIRE 4	Demand	0	45	30	0	66	52	0	20	20	70	0	0
	Production	24	32	32	32	32	32	32	32	32	32	-	-
	Inventory	24	11	13	45	11	-9	23	35	47	9	9	9
TIRE 5	Demand	0	0	25	0	18	0	86	12	51	0	0	27
	Production	-	13	16	16	29	32	32	32	32	16	16	-
	Inventory	0	13	4	20	31	63	9	29	10	26	42	15

In order to formulate a mathematical model for this production planning problem, we first define the sets, parameters and variables as follows:

Sets:

T : Set of different tires, index i,
H : Set of different heaters, index k,
P : Set of periods, index t,
TXH : Set of possible tire-heater combinations, index (i,k).

Parameters:

bc_i : backlog cost for tire i,
hc_i : holding cost for tire i,
nh_{kt} : number of heaters of type k in period t,
sc_{ik} : set up cost for tire-heater combination (i,k),
rt_{ik} : production rate per day for tire-heater combination (i,k),
rs_{ik} : production lost due to set up time for tire-heater combination (i,k),
rm_{ik} : number of running molds for tire-heater combination (i,k) at the beginning of the planning period,
nm_{ikt} : number of molds available for tire-heater combination (i,k) in period t,
dm_{it} : demand for tire i in period t.

Variables:

s_{it} : inventory of tire i in period t,

b_{it} : backlog of tire i in period t,

x_{ikt} : number of heaters of type k producing tire i in period t,

y_{ikt} : number of set ups for tire-heater combination (i,k) in period t.

The problem can then be formulated as the following integer program **(SOL)**:

$$\text{Min} \sum_{(i,k)\in TXH} \sum_{t\in P} sc_{ik}\, y_{ikt} + \sum_{i\in T} \sum_{t\in P} \left(bc_i\, b_{it} + hc_i\, s_{it}\right) \tag{1}$$

s.t.

$$b_{i,1} + \sum_{k|(i,k)\in TXH} (rt_{ik}\, x_{i,k,1} - rs_{ik}\, y_{i,k,1}) = dm_{i,1} + s_{i,1} \qquad \forall i\in T \tag{2}$$

$$s_{i,t-1} + b_{i,t} + \sum_{k|(i,k)\in TXH} (rt_{ik}\, x_{ikt} - rs_{ik}\, y_{ikt}) = dm_{it} + s_{it} + b_{i,t-1} \qquad \forall i\in T, \forall t\in P\setminus\{1\} \tag{3}$$

$$y_{i,k,1} \geq x_{i,k,1} - rm_{ik} \qquad \forall (i,k)\in TXH \tag{4}$$

$$y_{ikt} \geq x_{ikt} - x_{i,k,t-1} \qquad \forall (i,k)\in TXH,\ \forall t\in P\setminus\{1\} \tag{5}$$

$$\sum_{i|(i,k)\in TXH} x_{ikt} \leq nh_{kt} \qquad \forall k\in H, \forall t\in P \tag{6}$$

$$x_{ikt} \leq nm_{ikt} \qquad \forall (i,k)\in TXH, \forall t\in P \tag{7}$$

$$x_{ikt}, y_{ikt} \in \{0,1,2,3,\ldots\} \qquad \forall (i,k)\in TXH, \forall t\in P \tag{8}$$

The objective function (1) minimizes the total costs. The relevant costs are the set up cost, backlog cost and inventory holding cost. Constraints (2) and (3) are the demand constraints. The production lost due to set up times is properly taken into account. Constraints (4) and (5) are the set up forcing constraints. Further we have for each period a capacity constraint on the number of heaters (6) and on the number of molds (7). Finally we have the integrality constraints (8). Observe that the integrality constraint on the y_{ikt} variable is not strictly necessary.

2. Literature review

The problem we are facing is essentially a multi-item scheduling problem with multiple machine types, set up times, backlogging and a discrete production policy. Many papers have appeared in the literature describing scheduling problems complicated by one or more of these additional issues. To the best of our knowledge, this is the first paper that takes all of the additional issues simultaneously into account. Moreover, there is almost no literature describing a problem where a product can be produced on alternative non-identical machines. Our problem and solution method is related to the multi-item scheduling problem described by Lasdon and Terjung (1971) and the Discrete Lotsizing and Scheduling Problem with set up times studied by Cattrysse et al. (1993).

Lasdon and Terjung (1971) describe a multi-item scheduling problem at a major U.S. tire manufacturer. The relevant costs are the set up, backlog and holding costs. There is only one machine type available and no set up times are considered. In each period, each heater produces either nothing or at full capacity. Lasdon and Terjung solve the LP relaxation using a column generation scheme where for each item a subproblem is solved using a DP recursion. Eppen and Martin (1987) solve the Lasdon and Terjung

model, without backlogging but including the integrality constraints, using variable redefinition. They found that the relaxation of their formulation yields very tight bounds. Cattrysse et al. (1993) describe a heuristic for the Discrete Lotsizing and Scheduling Problem (DLSP) wit set up times. In the DLSP multiple items must be scheduled on a single machine. Backlogging is not allowed. In each period, at most one item can be produced and a set up is required if the production is changed to another item. The set up time is an integral number of periods, whereas in our problem set up times are a fraction of the unit period. DLSP assumes a discrete production policy. A column generation scheme is applied and the dual prices are calculated with a dual ascent method and subgradient optimization. The subproblems are solved using a DP recursion.
Other papers have researched set up times without the discrete production assumption. Manne (1958) formulates a multi-item lot sizing problem with one machine type. He incorporates set up times in his formulation and no backlogging is allowed. The multiple mode case is assumed, i.e. several items can be produced in the same time period. Manne explicitly treats all the possible Wagner-Whitin production schedules and proves that they are the only ones that need to be considered. In other words, he works with a full blown master. Trigeiro et al. (1989) also consider the Capacitated Lot Sizing Problem (CLSP) with set up times and the single machine, multiple mode assumption. The problem is solved by Lagrange relaxation where the dual prices are updated using subgradient optimization. The single item problems are solved by DP. They conclude that set up times make the problem much harder to solve. Diaby et al. (1992) also solved the CLSP with multiple resources using Lagrange duality. It was assumed that each item could be produced at the same production rate on each machine. Recently, Kang et al. (1999) extended the CLSP in two ways, by considering multiple parallel machines - with different production rates - and sequence-dependent set up costs. However, they do not take set up times into account. Vanderbeck (1998) studied lot sizing problems with start-up times. He considers the single-stage, single-mode and multi-item case with continuous set-ups, i.e. variable production is allowed, and sequence independent start-ups. He develops a branch-and-price scheme combining column generation and cutting planes. Further he examines complexity issues for different lot sizing models. Belvaux and Wolsey (1998) describe valid inequalities for different lot sizing models, including for multiple machines.

Applications of LP/IP techniques to production planning problems at tire manufacturers have already appeared in the literature. Lasdon and Terjung (1971) described the multi-item scheduling problem at Goodyear. Degraeve and Schrage (1997, 1998) report on a successful scheduling application at Bridgestone/Firestone Off-The-Road. The complex problem requires stacking of molds and assigning these stacks to heaters. The problem is solved using a column generation scheme. Reading this paper inspired the management of Solideal to come to us with their scheduling problem.

Our application differs in two important ways from the previous literature. First we consider simultaneously set up times and non-identical alternative machines. Second, the use of optimization based heuristics is directly motivated by industry requirements to generate several high quality solutions fast.

3. The algorithm

3.1 Decomposition

We decompose the problem into a master progam (SMAS) and a subproblem for every tire ($SSUB_i$).

We define a new set, variable and two parameter:

Q_i : set of columns generated for tire i,

z_{iq} : = 1 if column q is chosen for tire i, = 0 otherwise,

a_{iktq} : number of heaters active for tire-heater combination (i,k) in period t in schedule q,

$cost_{iq}$: cost of column q for tire i.

Master (SMAS):

$$\text{Min} \sum_{i \in T} \sum_{q \in Q_i} cost_{iq}\, z_{iq} \tag{9}$$

S.T.

dual price

$$\sum_{i|(i,k) \in TXH} \sum_{q \in Q_i} a_{iktq}\, z_{iq} \le nh_{kt} \qquad \forall t \in P,\ \forall k \in H \qquad \pi_{kt} \tag{10}$$

$$\sum_{q \in Q_i} z_{iq} = 1 \qquad \forall i \in T \qquad \mu_i \tag{11}$$

$$z_{iq} \in \{0,1\} \qquad \forall i \in T, \forall q \in Q_i \tag{12}$$

In the master problem, the objective function (9) minimizes the sum of the total costs. We retain the heater capacity constraint (10) as in the original model (SOL) and add a convexity constraint (11) for each tire. Finally we force the column variable z_{iq} to be binary (12). The column generation scheme allows us to solve an approximation of the problem SMAS where the integrality constraints (12) are relaxed. Because the subproblems are solved as an IP, the LP relaxation of SMAS will provide a tighter lower bound than the LP relaxation of the original problem SOL. Manne (1958) argues that this relaxation will be close to the IP optimum if the number of items is a lot larger than the number of periods. The solution to the overall problem (9-12) can be found by implementing a branch-and-price scheme (see e.g. Barnhart et al. 1998).

The information given by the dual prices π_{kt} and μ_i is used to generate new columns. For a specific tire type i and a set of dual prices we have the following single item subproblem:

Subproblem ($SSUB_i$): $\forall i \in T$

$$\text{Min} \sum_{k|(i,k) \in TXH} \sum_{t \in P} sc_{ik}\, y_{ikt} + \sum_{t \in P} (bc_i\, b_{it} + hc_i\, s_{it}) + \sum_{t \in P} \sum_{k|(i,k) \in TXH} \pi_{kt}\, x_{ikt} - \mu_i \tag{13}$$

S.T.

$$b_{i,1} + \sum_{k|(i,k) \in TXH} (rt_{ik}\, x_{i,k,1} - rs_{ik}\, y_{i,k,1}) = dm_{i,1} + s_{i,1} \tag{14}$$

$$s_{i,t-1} + b_{i,t} + \sum_{k|(i,k) \in TXH} (rt_{ik}\, x_{ikt} - rs_{ik}\, y_{ikt}) = dm_{it} + s_{it} + b_{i,t-1} \qquad \forall t \in P \setminus \{1\} \tag{15}$$

$$y_{i,k,1} \geq x_{i,k,1} - rm_{ik} \qquad \forall k \in H \mid (i,k) \in TXH \qquad (16)$$

$$y_{ikt} \geq x_{ikt} - x_{i,k,t-1} \qquad \forall k \in H \mid (i,k) \in TXH, \forall t \in P \setminus \{1\} \qquad (17)$$

$$x_{ikt} \leq nm_{ikt} \qquad \forall k \in H \mid (i,k) \in TXH, \forall t \in P \qquad (18)$$

$$x_{ikt}, y_{ikt} \in \{0,1,2,3,...\} \qquad \forall k \in H \mid (i,k) \in TXH, \forall t \in P \qquad (19)$$

The objective function (13) minimizes the reduced cost of the new column for tire i. A feasible production plan for an individual tire satisfies the following constraints: demand constraints (14 and 15), set up forcing constraints (16 and 17), capacity constraints on the number of molds (18) and integrality constraints (19).

The new columns which price out, i.e. have a negative reduced cost (13), are added to the master. The optimal solution to the subproblem $\left(x^*_{ikt}, y^*_{ikt}, b^*_{it}, s^*_{it}\right)$ defines the cost and production parameter of column q for tire i in the master as follows: $cost_{iq} = \sum_{k \mid (i,k) \in TXH} \sum_{t \in P} sc_{ik}\, y^*_{ikt} + \sum_{t \in P} \left(bc_i\, b^*_{it} + hc_i\, s^*_{it}\right)$ and $a_{iktq} = x^*_{ikt}$. Then the master (without the integrality constraint) is solved again and with the new dual prices we try to find new columns that price out. When no columns price out favorably, i.e. the reduced cost is nonnegative, we have found the optimal solution to the relaxed master (9-11).

3.2 Dynamic programming formulation for the subproblem

In order to get tight bounds, the subproblem must be solved including the integrality constraints (19). This single-item problem is a hard problem by itself. In our algorithm the subproblem $SSUB_i$ is solved using a dynamic programming recursion.

We first discuss the single item, single machine problem, where a tire can only be produced on one type of machine. The number of possible production plans in a p period problem equals (the number of molds + 1)^p. For 3 molds and 15 periods this amounts already to more than one billion production plans. The DP recursion has two state variables (Lasdon and Terjung 1971), namely the number of active molds (m) and the inventory position (unrestricted in sign) at the end of the period (inv). The stage t is the time period. The number of states for the inventory position of the tires with a single resource is limited in each period t as follows:

$$\forall i \in T, \exists! k \mid (i,k) \in TXH, \forall t \in P:$$

$$b_{it} \leq \sum_{j=1}^{t} dm_{ij} \qquad (20)$$

$$s_{it} \leq \min \begin{cases} \sum_{j=1}^{t} (rt_{ik}\, nm_{ikj} - rs_{ik}\, v_{ikj}) - \sum_{j=1}^{t} dm_{ij} \\ \sum_{j=t+1}^{p} dm_{ij} + rt_{ik} - 1 \end{cases} \qquad (21)$$

$$with\ v_{ik1} = \max\{0, nm_{ik1} - rm_{ik}\}, v_{ikj} = \max\{0, nm_{ikj} - nm_{i,k,j-1}\}$$

Of course we cannot have more backlog than the total demand over a period (20). For the stock we have two a priori upper bounds (21). First, we cannot have more stock than total production with the maximum number of molds minus total demand up to period t. Secondly, we do not want too much stock left at the end of the planning horizon, due to the positive inventory holding cost. However, because we adopt a discrete production policy, it is possible in an optimal plan to have some stock left at the end of the horizon, but not more than one day's production on one heater. The data that was given to us, allows us to model production and inventory as integer numbers. We also assume that there is no beginning inventory (see Section 4).
Starting from all the feasible states at stage $t-1$ and taking all the possible actions, i.e. deciding on the number of molds in period t, we can construct all the feasible states in stage t. If during this process we encounter a state that was already generated, we check if it can be reached with a lower cost. The state space in the DP algorithm can be reduced by introducing the following computational enhancements:

1. Dominance rule : if the inventory position in any stage violates the bounds of (20) or (21), then we can discard the corresponding node;
2. Upper bound 1: we first calculate the cost if we would produce with the maximum number of molds available. This is an upper bound on the total cost and any node at any stage exceeding this cost can be discarded;
3. Upper bound 2: we are only interested in those proposals with a negative reduced cost. Hence we can delete nodes at any stage if they have a non-negative reduced cost.

The same reasoning can be extended to the single item, multiple machine problem, where a tire can be produced both on a regular heater and a BOM. The DP recursion now has three state variables: the number of molds active on the regular heaters (m_1), the number of molds active on the BOM heaters (m_2) and the inventory position at the end of the period (inv).

A limited state space is an important condition for the efficiency of this DP algorithm. In practice the number of molds is limited to an order of 10 and the demand is limited to an order of 1500. At Solideal, the tires for which this demand is exceeded are called 'fast movers'. They are produced continuously with dedicated production equipment. As such, the total production capacity of this dedicated equipment can be subtracted from the demand in a preprocessing step of our algorithm (see Section 4). The remaining peak demand, typically below 1500, remains to be scheduled.

3.3 Lagrange relaxation

The master problem (SMAS) is degenerate and hence does not converge rapidly if we use a standard simplex algorithm to determine the dual prices. The explicit requirement to generate high quality solutions in a short computational time promted us to look for faster ways to calculate the LP relaxation. Lagrange relaxation can be used as an alternative way of finding the dual prices and counters the problem of degeneracy. Lagrange relaxation has the additional advantage that during the subgradient phase feasible solutions are generated. As such it meets the company requirement that various alternative production proposals should be presented. In their paper on the discrete lotsizing and scheduling problem, Cattrysse et al. (1993) also use this technique to solve a degenarate master. Fisher (1985) gives a good and practical

introduction to Lagrange relaxation. We relax the heater capacity constraints (10) in SMAS to obtain the Lagrange problem. Then we use subgradient optimization to compute approximately the dual prices. These dual prices are the input for the subproblems.

The Lagrange dual problem of the master becomes **(SLAG)**:

$$\text{Min} \sum_{i\in T}\sum_{q\in Q_i}(cost_{iq} + \sum_{t\in P}\sum_{k|(i,k)\in TXH} a_{iktq}\,\pi_{kt})\,z_{iq} - \sum_{t\in P}\sum_{k\in H}\pi_{kt}\,nh_{kt} \qquad (22)$$

S.T.

$$\sum_{q\in Q_i} z_{iq} = 1 \qquad \forall i\in T \qquad \mu_i \qquad (23)$$

$$z_{iq}\in\{0,1\} \qquad \forall i\in T, \forall q\in Q_i \qquad (24)$$

We have implemented a subgradient optimization scheme for setting the dual prices π_{kt}. During a fixed number of iterations, the dual prices are updated:

$$\pi_{kt}^{\,r+1} = \max\left\{0,\ \pi_{kt}^{\,r} - \lambda\left(nh_{kt} - \sum_{i|(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\,z_{iq}^{\,r}\right)\right\} \qquad \forall t\in P,\ \forall k\in H \qquad (25)$$

with

$$\lambda = \frac{\omega\left(UB - Z_{SLAG}(\pi_{kt}^{\,r})\right)}{\sum_{t\in P}\sum_{k\in H}\left(nh_{kt} - \sum_{i|(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\,z_{iq}^{\,r}\right)^2} \qquad (26)$$

$Z_{SLAG}(\pi_{kt}^{\,r})$ is the optimal objective value (22) of the Lagrange dual problem for a given set of dual prices $\pi_{kt}^{\,r}$.The Lagrange dual is easily solved by taking for each tire the column with the lowest cost. UB is an upper bound of the full blown master. This is the best known feasible IP solution. ω is initially fixed at 1.5 and is halved every time when the lower bound has failed to increase for a fixed number of iterations. In order to generate good columns, we also need the dual prices μ_i of the convexity constraints (23). The dual price of the convexity constraint for tire i equals the cost of the optimal column, as defined in the first term of (22). At every iteration of the subgradient optimization, we test whether the Lagrange solution is primal feasible and if required we adapt the upper bound. The initial dual prices π_{kt} are set so that the dual objective value equals an initial primal upper bound.

3.4 Upper bounds

Because backlogging is allowed, every production plan which does not violate constraint (7) is feasible with respect to constraints (2)-(5) and we only have to take care that the total heater capacity is not exceeded. This allows us to easily construct initial upper bounds. A simple heuristic works as follows. Begin with the first tire and schedule it optimally, using the DP recursion. Adapt the heater and mold capacities and schedule now the second tire optimally, taking into account the new capacities. We repeat this procedure until all the tires are scheduled. To obtain better upper bounds, we repeat this heuristic, but we schedule the tires in a different order. We use

some different ranking principles. First we rank the tires by increasing and decreasing total backlog. Total backlog is here defined as $\sum_{t \in P} dm_{it}\,(p+1-t)$, with p the number of periods. A second ranking principle is the total demand. Here we also rank first by increasing and then by decreasing total demand. For each of these cases we have also implemented a heuristic where we first rank the tires which can be produced on one machine only, after that we rank the tires which can be produced on two machines. Finally we also schedule the tires in the order as they appear in the input file. This heuristic is used to construct initial upper bounds and initial columns at the first step of our algorithm. The heuristic is very fast, but does not produce the high quality proposals which management wants. The schedules can be improved manually by simple visual inspection of the planners. This prompted us to search for optimization based heuristics which resulted in the column generation based heuristic.

At the end of the column generation process, when we have found the LP relaxation, we try to construct a good feasible IP solution, using an LP fixing heuristic as follows. We first round down the LP solution. Then in iterative steps, we try to increase production. We first look at the variables with a fractional part greater than 0.6, next come the variables with a fractional part greater than 0.3 and finally the remaining variables are investigated. Within these loops, we investigate the tires in the same order as with the initial heuristic. If we have a fractional LP solution for a certain x_{ikt} variable, we increase the production to the next integer if the capacity constraint is not violated. We adapt the capacities and repeat this procedure for another x_{ikt} variable. If after this procedure there are still some tires with backlog and there is still capacity available, we try to increase production for the variables which were not fractional. If we find columns which were not already added, we also put them in the master. The master is then submitted to the branch-and-bound solver of LINDO with the best upper bound given as described above.

3.5 Lower bound calculation

During the column generation process the value of the objective function (9) provides a primal upper bound on the LP relaxation of SMAS. We can also compute a dual lower bound on the master (Lasdon and Terjung 1971, Wolsey 1998 or Martin 1999). Let z^t be the objective value of SMAS at iteration t and let rc_i^t be the reduced costs of columns which we have added at iteration t. If no column was added for an item, the reduced cost equals zero. The lower bound on the master is than as follows:

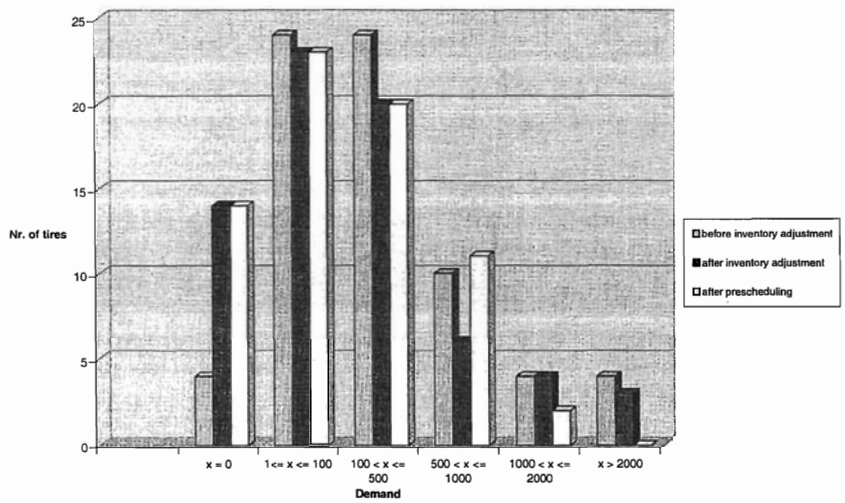
$$\text{LB} = z_t + \sum_{i \in T} rc_i^t \tag{27}$$

Based on this lower bound, we implement a stopping criterion. If the percentage difference between the current master solution and the lower bound is less than a specific value, e.g. 0.5%, we stop. The optimal dual prices are needed to calculate this lower bound, therefore we cannot apply it if we solve the master with Lagrange relaxation.

4. The issue of fast movers and other industrial complications

In practice, the company disposes of several dedicated lines of heaters for the fast mover tires. These are tires with a high demand, typically higher than 1500 over a 30 day time horizon. To satisfy demand, these tires have to be produced continuously. This means that we have to adjust the demand of these so-called fast movers and also the number of molds and heaters available to reflect the production on the dedicated equipment. We now only have to solve the competition of the remaining demand for the remaining resources. In practice the data will be sufficiently reduced in order to efficiently use our algorithm. In Figure 1 we can see an example of the effect of this 'prescheduling' on a large set of tires. On the x-axis we see different ranges for the demand and the y-axis shows the number of tires within that demand category. We observe that after the prescheduling the number of tires with a high demand has decreased considerably, while the remaining demand of the fast movers is added to a lower demand category.

Figure 1 The effect of inventory adjustment and prescheduling on demand.



During discussion with the planners, the problem was raised that there might be too many set ups in the morning shift. The number of mold changes per day should be limited because such a mold change requires extensive use of limited manpower and equipment. From our results it has not yet been apparent that this would be a major issue. However, this industrial complication could be implemented in the following way (see also Lasdon and Terjung 1971):

Define:

w_{itq} : number of mold changes for tire i in period t in schedule q,

c_t : maximum number of mold changes allowed in period t.

The new constraint then becomes:

$$\sum_{i\in T}\sum_{q\in Q_i} w_{iiq}\, z_{iq} \le c_t \qquad \forall t \in P \qquad \tau_t \qquad (28)$$

The objective function of the subproblem becomes:

$$\sum_{k|(i,k)\in TXH}\sum_{t\in P}(sc_{ik}+\tau_t)\, y_{ikt} + \sum_{t\in P}(bc_i\, b_{it} + hc_i\, s_{it}) + \sum_{t\in P}\sum_{k|(i,k)\in TXH}\pi_{kt}\, x_{ikt} - \mu_i \qquad (29)$$

The Lagrange relaxation algorithm has to be adapted too. Dualizing constraint (27) gives the following objective function for the Lagrange dual problem:

$$\sum_{i\in T}\sum_{q\in Q_i}(cost_{iq} + \sum_{t\in P}\sum_{k|(i,k)\in TXH} a_{ikiq}\,\pi_{kt} + \sum_{t\in P} w_{iiq}\,\tau_t)\, z_{iq} - \sum_{t\in P}\sum_{k\in H}\pi_{kt}\, nh_{kt} - \sum_{t\in P}\tau_t\, c_t \qquad (30)$$

Another problem arises when some molds can be used both for a regular heater and a BOM heater. This is currently not yet the case. However, in the foreseeable future, given plant capacity expansion, this might potentially occur. We can accommodate this request by adapting the DP recursion for the multiple machine problem. Assume that we have for a specific tire r_1 regular molds, r_2 BOM molds and r_3 molds which can be used on both heaters. In the DP recursion, the state variables should then be limited as follows:

$$m_1 \le r_1 + r_3$$
$$m_2 \le r_2 + r_3$$
$$m_1 + m_2 \le r_1 + r_2 + r_3$$

where m_1 is the number of active molds on the regular heaters and m_2 is the number of active molds on the BOM heaters.

Taking into account some other practical considerations, it is possible to further limit the state space of the DP algorithm. For example, management could state that the backlog may never be larger than half of the total demand over the scheduling horizon. Such a constraint is easily incorporated in the DP recursion. In this way, the state space is further reduced and the algorithm will be speeded up. We must however remark that the addition of such managerial concerns can lead to an infeasible problem. We could also initially use this rule and when no more columns price out, we try again to generate columns without restriction on the backlog.

5. Computational results

The data which we used to test the algorithms consist of a real life order book. We used the demand figures which we obtained after preprocessing the data, as described before. For our experiment, we considered three different time horizons: 10, 20 and 30 periods and we derived three tire sets with 10, 20 and 30 tires. Combined they give nine data sets with different time horizons and number of tires. The data sets are abbreviated as T*a_b*, where *a* refers to the number of items and *b* refers to the number of periods. The production data consist of the following information for every tire: we indicate in which heater(s) it can be cured, the number of available molds, the daily production rate and the set up loss in units of production. Some tires can be cured both on a regular heater and a BOM heater.

The choice of the problem sizes was motivated by the daily practice of the company. The large time horizon - up to 30 days - is inherently linked to the manufacturing environment with large series production. The tire data sets also have typical

industrial sizes. The chairman was first surprised that we only scheduled 30 tires, given the product range of 475 tires. After drawing his attention to our data analysis, indicating that at most 30 tires are competing for the same resources and as such must be scheduled simultaneously, he acknowledged that this was in fact the case. The smaller data sets with 10 time periods, which still could be classified as medium sized, have less industrial relevance but are used here to show the quality of our algorithms.

All our experiments were performed on a modest 233 mHz pentium II computer with a 64 MB memory. The algorithms are programmed in Fortran 77 using the WATCOM Fortran compiler version 10.6 and linked with the industrial LINDO optimization library version 5.3 (Schrage 1995). All the computation times are given in hours:minutes:seconds.

Trigeiro et al. (1989) state that the tightness of the capacity constraint is a good measure for the difficulty of the problem with set up times. However, in the presence of set up times and multiple resources, measuring the a priori tightness of the capacity constraint is impossible. Determining a priori the capacity needed requires a decision concerning the production and number of set ups (Trigeiro et al. 1989) and a rule for splitting the demand over the different machines. Kang et al. (1999) remark that the use of multiple machines with different production times makes capacity utilization dependent on the solution. In our experiments, we will measure the tightness of the capacity constraint by calculating the capacity utilization for the best IP solution found at the end of our procedure.

In a first experiment we have solved the problems with the decomposition method, where the master is solved using the LP solver of LINDO. We stopped if we could not add any column or if the relative difference between the objective value of the restricted master and the lower bound (27) was less than 0.0005%. We will further refer to this algorithm as M1. These results are given in Table 4. The first column indicates the problem set. The second column, labeled 'Capacity', gives the number of heaters available for the three types. In the third column the tightness of the capacity constraints for the three heaters, calculated as percentage capacity usage, is given. The following columns give the cost of the best initial heuristic solution, the gap compared to the LP relaxation and the time to calculate it. The column labeled 'LP' gives the LP relaxation of the problem. If there was still a difference between this LP solution and the lower bound, we reported the lower bound. In the next two columns we give the time needed to calculate the LP solution and the number of columns which were added. Further, we give the solution found after applying the fixing heuristic and the gap with the LP relaxation. The time to calculate this upper bound was typically less than half a second. Finally, we report for benchmarking purposes only, the best IP solution found after applying a branch-and-bound algorithm, the gap with the LP solution and the time. In general we see that for the high capacitated problems more columns are added. Problems with a loose capacity constraint have generally a smaller IP gap and the initial heuristic gives already a good solution. We observe that especially data sets T10_20 and T10_30 are very degenerate, resulting in a large CPU time for finding the LP and a large number of columns added. Finally note that all the problems of the smallest data set T10_10 could be solved to optimality within a short CPU time.

In a second experiment we tried to counter the problem of degeneracy by imposing a less severe stopping criterion. Now we stop if the gap between the LP relaxation of the

restricted master and the lower bound is less than 0.5 %. This algorithm is further called M2. In Table 5 we give the lower bound on the master, as calculated in (27), and the LP solution of the restricted master in the third and fourth column. The gaps are calculated relative to this LP lower bound. From the results in Table 5, we see that less columns are added and the LP computation time is usually shorter. The IP solutions found are usually worse, but sometimes also better.

In a third experiment we used Lagrange relaxation to solve the master. This algorithm is referred to as M3. Because the Lagrange solution is only an approximation of the optimal dual variables, it is possible that the algorithm will terminate even if not all the columns which are necessary to obtain the true LP relaxation value are added. In practice, the value of the restricted master at termination of the algorithm will be close to the true LP relaxation. In theory however, we cannot rely on this bound. In Table 6 we find some new columns. 'Lagrange LB' gives the value of the master at termination as calculated with subgradient optimization. 'LP Master' then gives the true value of the restricted master at termination. This value is calculated with LINDO. The column labeled 'Lagrange Heuristic' gives the value of the best IP solution at the end of the column generation process as found during the subgradient optimization. Such a solution is not always found. Next we give the best solution of the fixing heuristic. The next column calculates the gap between the best heuristic solution and the LP relaxation, as found with M1, and gives as such an indication of the quality of the heuristic solution. The one but last column of Table 6 gives an indication of the quality of the best solution that exists in the restricted master, after we have allowed a LINDO branch-and-bound search for 9,999,999 pivots. The last column gives the corresponding computation times, which are very large if the optimum solution is not found within this limit. However, if we lower the pivot limit to 100,000 only 5 out of the 60 test instances will have a slightly worse IP solution. The corresponding computation times, however, are drastically reduced for the difficult instances. This indicates that it might be worthwhile to do some extra pivots after a fixing heuristic, but not too many.

Table 7 compares for all the data sets the average LP time and the average number of columns added for the three algorithms. For M2 and M3 we also give the average improvement in percentage, compared to M1. We see that M3, and to a lesser extent M2, is very successful in reducing the average LP time. The percentage improvement varies approximately between 40 and 90 percent. The average number of columns has decreased with 30 to 80 percent. Especially for the two most degenerate problems, T10_20 and T10_30, we observe a substantial improvement. As our heuristic solutions are found at approximately the LP time, the average IP time is a less accurate measure for the performance of our algorithms because it actually measures the performance of the LINDO branch-and-bound algorithm. Finally in Table 8 we compare the best heuristic and IP solutions for the three algorithms. We calculate the average gap, compared with the LP solution found with M1. Normally, this exact LP bound is not available with algorithms M2 and M3 and the only purpose of this calculation is to compare the solutions. We see that in general the gaps are very small. On average, M2 and M3 have the same gap for the best heuristic solution (1.6%) and also for the IP solution (approximately 1%). M1 has a worse heuristic gap (2.2%), but a better IP gap (0.8%). It appears that the large reduction in columns resulting from the Lagrange relaxation, does not substantially influence the quality of the IP solution.

TABLE 4 Solution with M1

Problem	Capacity	Capacity Tightness	Initial Heuristic	Gap (%)	Time Heuristic	LP	Time LP	Number Columns	Fixing Heuristic	Gap (%)	IP	Gap(%)	Time IP
T10_10	3/2/2	100 / 100 / 90	242,983	14.43	1	212,346	21	253	212,346	0.00	212,346	0.00	21
	3/3/3	100 / 90 / 73	205,676	10.20	2	186,640	10	195	186,640	0.00	186,640	0.00	10
	4/4/4	85 / 70 / 55	147,583	2.18	1	144,440	2	72	144,440	0.00	144,440	0.00	2
	5/5/5	68 / 60 / 40	136,477	0.48	2	135,819	2	42	135,819	0.00	135,819	0.00	3
	6/6/6	57 / 53 / 32	132,159	0.03	2	132,122	2	15	132,122	0.00	132,122	0.00	2
	7/7/7	49 / 46/ 27	131,636	0.00	2	131,636	2	14	131,636	0.00	131,636	0.00	2
T10_20	5/4/4	100 / 100 / 100	943,820	27.74	37	738,839°	34:25	1,224	740,193	0.18	740,193	0.18	34:29
	5/5/5	100 / 88 / 91	672,853	6.77	31	630,190	17:22	886	631,542	0.21	631,542	0.21	17:49
	6/5/5	100 / 88 / 91	487,110	8.07	32	450,726	54:59	1,478	(543,938)	20.68	452,037	0.29	55:41
	6/6/6	100 / 78 / 73	479,756	7.45	1:20	446,486.9°	1:13:54	1,242	479,756	7.45	447,891	0.31	1:15:30
	7/7/7	89 / 64 / 64	427,444	0.05	50	427,210	2:29	73	427,210	0.00	427,210	0.00	2:29
T10_30	4/3/3	100 / 100 / 100	2,692,838	22.78	31	2,193,135.5°	2:09:52	2,207	2,195,462	0.11	2,195,462	0.11	6:40:33*
	5/3/3	100 / 100 / 100	2,113,022	28.03	30	1,650,353	2:50:16	2,331	1,659,367	0.55	1,659,367	0.55	8:28:58*
	5/4/4	98 / 83 / 83	1,099,028	23.10	1:11	892,757.5	3:06:01	2,918	952,373	6.68	893,424	0.07	3:06:39
	5/5/5	98 / 67 / 65	815,908	7.37	55	759,907	39:00	1,424	759,907	0.00	759,907	0.00	39:07
	7/5/5	75 / 61 / 63	448,597	0.40	1:06	446,810.3	10:31	459	448,329	0.34	446,894	0.02	10:39
	7/7/7	72 / 48 / 45	442,426	0.05	1:18	442,216.3	2:13	52	442,300	0.02	442,254	0.01	2:13
T20_10	5/5/5	96 / 100 / 98	282,652	32.85	10	212,756.4°	4:37	363	227,410	6.89	220,134	3.47	4:55
	6/6/6	88 / 85 / 87	193,987	28.38	24	151,103	1:03	162	151,103	0.00	151,103	0.00	1:03
	7/7/7	76 / 67 / 81	156,637	10.17	13	142,171.5	21	88	142,259	0.06	142,231	0.04	21
	8/8/8	66 / 63 / 68	145,029	3.15	26	140,599.3	29	42	141,114	0.37	140,658	0.04	29
T20_20	7/7/7	100 / 100 / 99	1,683,965	30.52	1:33	1,290,157.1°	27:40	1,219	1,298,508	0.65	1,290,899	0.06	28:14
	8/8/8	100 / 86 / 88	1,377,183	28.82	1:01	1,069,071	6:29	586	1,091,571	2.10	1,070,088	0.10	6:36
	10/8/8	100 / 86 / 88	812,720	17.89	1:00	689,380.6°	6:14	554	696,917	1.09	690,085	0.10	6:23
	11/8/8	100 / 86 / 88	655,169	17.95	1:08	555,462	5:30	465	565,552	1.82	556,155	0.12	5:34
	12/8/8	95 / 86/ 88	527,707	1.86	54	518,064	3:53	312	525,383	1.41	518,272	0.04	3:57
	9/9/9	100 / 76 / 78	1,078,645	23.70	1:26	872,007.3°	5:00	411	874,690	0.31	873,136	0.13	5:03
	10/10/10	100 / 77 / 63	805,487	16.90	2:06	689,030°	6:24	429	702,191	1.91	689,869	0.12	6:29
	11/11/11	100 / 73 / 54	647,537	16.66	1:49	555,082.4°	5:55	380	561,117	1.09	555,105	0.00	5:56
	12/12/12	95 / 68 / 49	520,831	0.61	2:29	517,684.5	4:34	198	(521,911)	0.82	517,837	0.03	4:35

Problem	Capacity	Capacity Tightness	Initial Heuristic	Gap (%)	Time Heuristic	LP	Time LP	Number Columns	Fixing Heuristic	Gap (%)	IP	Gap (%)	Time IP
T20_30	10/7/7	92 / 89 / 70	1,320,145	45.88	1:04	904,931.5	36:13	1,354	945,496	4.48	911,613	0.74	1:12:16
	11/7/7	83 / 89 / 71	1,004,571	45.96	1:37	688,243.4	32:44	1,148	803,719	16.78	696,058	1.14	3:27:50
	8/8/8	100 / 76 / 62	2,130,723	18.75	1:16	1,794,302°	12:03	790	1,796,592	0.13	1,794,521	0.01	12:18
	9/7/6	100 / 93 / 78	2,016,062	46.71	1:17	1,374,212.5°	1:10:31	1,664	1,415,189	2.98	1,415,189	2.98	4:40:22*
	9/9/9	99 / 70 / 54	1,570,003	24.57	1:38	1,260,326.2	12:56	692	1,279,361	1.51	1,265,617	0.42	19:24
	10/10/10	92 / 63 / 47	1,110,079	25.01	2:05	887,991	16:35	687	887,991	0.00	887,991	0.00	16:37
	11/11/11	85 / 58 / 42	777,242	15.78	2:54	671,297.9	8:34	423	735,835	9.61	671,974	0.10	8:38
	12/12/12	78 / 53 / 38	631,486	0.51	2:58	628,282	5:50	215	(634,389)	0.97	628,457	0.03	5:52
T30_10	5/5/5	98 / 100 / 98	331,800	31.07	10	253,150.5	3:44	327	265,475	4.87	260,615	2.95	4:08
	6/5/5	97 / 100 / 100	294,829	35.90	9	216,939.3	4:32	405	230,019	6.03	218,849	0.88	4:35
	7/5/5	89 /100 / 98	273,430	43.07	9	191,120	7:08	454	200,639	4.98	197,607	3.39	8:53
	8/5/5	78 / 100 / 100	255,184	39.74	16	182,619	5:45	389	191,814	5.04	188,261	3.09	5:51
	9/5/5	69 / 100 / 98	243,929	34.31	11	181,611	6:56	373	190,806	5.06	188,098	3.57	7:10
	6/6/6	97 / 85 / 87	224,410	22.97	11	182,487.3	29	217	191,768	5.09	182,554	0.04	29
	7/7/7	89 / 67 / 81	185,759	21.41	15	153,001.5	21	130	153,089	0.06	153,089	0.06	21
	9/6/6	69 / 85 / 87	173,510	17.91	20	147,159	1:04	167	147,159	0.00	147,159	0.00	1:04
	9/6/7	69 / 77 / 81	164,286	13.06	11	145,304	18	94	145,304	0.00	145,304	0.00	18
	9/7/7	69 / 67 / 81	154,243	7.49	14	143,492.5	20	92	143,580	0.06	143,580	0.06	20
T30_20	8/5/5	100 / 100 / 100	3,110,921	30.27	42	2,388,004	48:19	1,250	2,394,480	0.27	2,388,291	0.01	48:24
	10/10/10	100 / 78 / 82	1,436,947	28.38	1:26	1,119,310.8	4:02	453	1,127,514	0.73	1,121,383	0.19	4:06
	14/8/8	100 / 99 / 100	843,870	47.97	1:35	570,284.1	32:42	1,198	580,465	1.79	573,835	0.62	2:56:52
	15/8/8	100 / 98 / 98	669,224	24.51	1:24	537,500.5°	24:50	1,018	547,309	1.82	543,530	1.12	3:08:41*
	16/8/8	97 / 95 / 100	558,111	4.66	1:18	533,268.8	14:11	765	555,115	4.10	555,115	4.10	2:37:44*
	16/9/9	92 / 91 / 88	534,738	0.78	1:40	530,576.7	3:47	207	534,659	0.77	530,612	0.01	3:48
	17/10/10	88 / 78 / 79	530,151	0.08	1:52	529,709.7	2:36	103	(530,207)	0.09	529,711	0.00	2:36
T30_30	14/8/8	87 / 88 / 77	1,141,096	60.92	1:37	709,109°	38:15	1,303	770,829	8.70	770,829	8.70	3:56:08*
	15/8/8	81 / 88 / 77	870,039	32.52	1:44	656,555.8°	30:55	1,050	687,316	4.69	687,316	4.69	3:37:12*
	16/8/8	79 / 82 / 75	715,809	10.24	1:29	649,339.7	15:08	706	659,206	1.52	659,206	1.52	3:12:11*
	16/9/8	79 / 77 / 69	656,919	1.48	2:06	647,332.4°	7:11	353	(657,540)	1.58	648,198	0.13	9:05
	16/9/9	76 / 75 / 67	650,472	0.63	1:53	646,376.8	5:32	248	(651,644)	0.81	646,470	0.01	5:35

° : The lower bound on the master is given if it deviates from the LP relaxation, This deviation. however is less than 0.0005 %,

* : Pivot limit of 9,999,999 was exceeded

() : LP rounding heuristic is worse than initial heuristic

TABLE 5 Solution with M2

Problem	Capacity	Master LB	Master LP	Time LP	Number Columns	Fixing Heuristic	Gap (%)	IP	Gap (%)	Time IP
T10_10	3/2/2	211,984	212,860	14	195	212,860	0.41	212,860	0.41	14
	3/3/3	185,749.1	186,640	5	157	186,640	0.48	186,640	0.48	5
	4/4/4	144,356.7	144,801.5	2	71	144,440	0.06	144,440	0.06	3
	5/5/5	135,698	135,819	2	33	135,819	0.09	135,819	0.09	2
	6/6/6	131,920.5	132,159	0	14	132,159	0.18	132,159	0.18	1
	7/7/7	131,636	131,636	2	14	131,636	0	131,636	0	2
T10_20	5/4/4	736,008.8	739,617.2	23:04	1008	740,193	0.57	740,193	0.57	23:05
	5/5/5	627,901.7	630,822.8	8:42	653	631,542	0.58	631,542	0.58	8:44
	6/5/5/	450,101.6	452,208.8	7:53	612	455,503	1.20	455,310	1.16	42:38
	6/6/6	445,735.3	446,952.2	9:53	454	446,785	0.24	446,785	0.24	9:54
	7/7/7	425,584.3	427,444	1:44	44	427,444	0.44	427,444	0.44	1:44
T10_30	4/3/3	2,185,783.5	2,196,188.8	1:22:43	1837	2,218,865	1.51	2,218,865	1.51	5:55:48*
	5/3/3	1,645,242.5	1,652,757.8	1:21:03	1825	1,661,848	1.01	1,661,848	1.01	6:15:49*
	5/4/4	889,845.2	893,817.3	42:47	1292	901,506	1.31	895,716	0.66	55:42
	5/5/5	757,068.3	760,839.7	17:27	1052	765,311	1.09	765,311	1.09	3:07:42*
	7/5/5	446,050.1	447,670.1	4:08	208	448,537	0.56	448,376	0.52	4:18
	7/7/7	441,865	442,254	1:16	24	442,254	0.09	442,254	0.09	1:16
T20_10	5/5/5	211,768.2	212,822.5	3:32	332	224,514	6.02	219,106	3.47	3:39
	6/6/6	150,688.4	151,397	36	137	151,397	0.47	151,397	0.47	36
	7/7/7	141,776.6	142,432	26	79	142,432	0.46	142,432	0.46	26
	8/8/8	140,594.8	141,092.4	17	36	141,533	0.67	141,142	0.39	17
T20_20	7/7/7	1,285,911	1,291,825.2	17:40	977	1,297,344	0.89	1,296,594	0.83	1:45:36*
	8/8/8	1,064,580.1	1,069,738.4	5:28	555	1,090,639	2.45	1,071,486	0.65	5:41
	10/8/8	687,432.5	690,818.3	4:56	511	703,637	2.36	691,788	0.63	5:28
	11/8/8	553,965.5	555,899.5	4:16	419	564,671	1.93	557,023	0.55	4:21
	12/8/8	517,208.2	519,193.2	3:03	256	521,397	0.81	518,267	0.20	3:05
	9/9/9	870,207.6	872,664.6	4:20	385	873,636	0.39	873,038	0.33	4:24
	10/10/10	686,802.7	689,178.8	4:41	388	692,756	0.87	692,756	0.87	4:51
	11/11/11	553,144.2	555,355.6	4:44	336	560,715	1.37	556,935	0.69	4:45
	12/12/12	516,733.5	518,446.2	3:21	145	518,691	0.38	518,257	0.29	3:21

Problem	Capacity	Master LB	Master LP	Time LP	Number Columns	Fixing Heuristic	Gap (%)	IP	Gap (%)	Time IP
T20_30	10/7/7	901,788.4	906,283.5	23:43	1111	935,312	3.72	923,214	2.38	3:50:23*
	11/7/7	686,168.6	689,360.5	24:24	1057	802,627	16.97	802,627	16.97	3:27:50*
	8/8/8	1,788,194.8	1,796,464.4	8:05	689	1,806,114	1.00	1,804,918	0.94	8:32
	9/7/6	1,370,249.1	1,376,078.4	45:23	1473	1,465,175	6.93	1,465,175	6.93	4:07:39*
	9/9/9	1,257,068.5	1,263,166.8	7:18	582	1,273,281	1.29	1,270,889	1.10	17:01
	10/10/10	886,988.3	889,502.1	8:05	513	928,032	4.63	891,560	0.52	8:24
	11/11/11	669,756.8	672,192.4	6:59	363	685,761	2.39	674,388	0.69	7:03
	12/12/12	626,894.7	628,417	4:52	172	628,847	0.31	628,457	0.25	4:53
T30_10	5/5/5	252,660.9	253,424.3	2:34	304	259,170	2.58	255,933	1.30	2:35
	6/5/5	216,186.2	217,153.2	3:01	358	229,653	6.23	220,213	1.86	3:03
	7/5/5	190,499.8	191,301.9	4:25	379	193,824	1.74	193,728	1.69	7:41
	8/5/5	181,992.6	182,851.9	3:21	320	185,319	1.83	185,166	1.74	3:22
	9/5/5	181,167.8	181,792.4	5:25	300	184,315	1.74	184,219	1.68	4:29
	6/6/6	182,035.9	182,781.3	36	214	186,560	2.49	182,848	0.45	36
	7/7/7	152,790.9	153,275.1	18	134	154,836	1.34	153,337	0.36	18
	9/6/6	146,508.6	147,159	45	139	147,159	0.44	147,159	0.44	46
	9/6/7	145,187.8	145,672.3	17	97	145,330	0.10	145,330	0.10	17
	9/7/7	143,103.8	143,753	17	78	143,753	0.45	143,753	0.45	18
T30_20	8/5/5	2,383,551	2,393,775	30:04	1088	2,402,852	0.81	2,399,648	0.68	30:23
	10/10/10	1,117,607	1,121,771.9	3:37	433	1,130,238	1.13	1,122,696	0.46	3:38
	14/8/8	568,236.9	570,899.8	20:39	1066	580,686	2.19	577,798	1.68	1:30:32*
	15/8/8	535,442.8	537,952.9	16:21	894	549,876	2.70	549,876	2.70	2:52:23*
	16/8/8	531,860.7	533,847.1	8:27	613	539,591	1.45	534,822	0.56	9:23
	16/9/9	527,962	531,255	1:58	162	(536,396)	1.60	531,591	0.69	2:00
	17/10/10	529,106.6	529,778.7	1:48	84	(530,513)	0.27	530,032	0.17	1:48
T30_30	14/8/8	706,276.8	709,697.7	28:26	1195	769,070	8.89	716,179	1.40	31:28
	15/8/8	654,290.4	657,378.6	21:36	950	678,225	3.66	678,225	3.66	3:18:28*
	16/8/8	648,040.8	649,868.2	11:28	674	677,372	4.53	677,372	4.53	3:14:40*
	16/9/8	645,834.1	647,961.1	3:29	283	(656,949)	1.72	649,184	0.52	3:44
	16/9/9	644,972.4	646,595.3	2:30	191	(663,625)	2.89	647,433	0.38	2:34

* : Pivot limit of 9,999,999 was exceeded

() : LP rounding heuristic is worse than initial heuristic

TABLE 6 Solution with M3

Problem	Capacity	Lagrange LB	LP Master	Time LP	Number Columns	Lagrange Heuristic	Fixing Heuristic	Gap (%)	IP	Time IP
T10_10	3/2/2	212,345.9	212,346	9	109	228,928	212,346	0.00	212,346	9
	3/3/3	186,639.9	186,640	4	81	197,662	186,640	0.00	186,640	4
	4/4/4	144,439.9	144,440	2	41	144,440	144,440	0.00	144,440	2
	5/5/5	135,781.2	135,819	2	23	135,831	135,819	0.00	135,819	3
	6/6/6	132,072.7	132,122	2	17	132,134	132,122	0.00	132,122	2
	7/7/7	131,623.9	131,636	2	13	-	131,636	0.00	131,636	2
T10_20	5/4/4	738,840.7	738.918	10:41	408	849,937	740,193	0.18	740,193	10:42
	5/5/5	630,164.1	630,191.3	3:02	207	670,698	631,550	0.22	631,550	3:02
	6/5/5/	450,726	450,726	2:57	216	452,658	(453,938)	0.43	452,225	2:59
	6/6/6	446,487	446,487	3:01	125	447,891	(449,699)	0.31	447,891	3:02
	7/7/7	427,210	427,210	1:37	40	427,210	427,210	0.00	427,210	1:37
T10_30	4/3/3	2,193,218.8	2,193,582	33:15	770	2,608,765	2,198,774	0.26	2,198,774	5:29:31*
	5/3/3	1,650,664.1	1,650,905.6	32:08	819	-	1,670,550	1.22	1,670,550	6:43:20*
	5/4/4	892,781.5	892,781.5	15:44	537	1,069,840	896,684	0.44	893,392	16:14
	5/5/5	759,907	759,907	4:36	233	813,780	759,907	0.00	759,907	4:37
	7/5/5	446,810.4	446,810.4	2:15	86	448,597	448,329	0.34	446,894	2:17
	7/7/7	442,100.9	442,254	1:38	40	442,254	442,254	0.01	442,254	1:38
T20_10	5/5/5	212,677.4	212,777.2	1:44	221	-	222,767	4.71	217,977	1:45
	6/6/6	151,103	151,103	29	108	153,332	151,103	0.00	151,103	30
	7/7/7	142,171.5	142,171.5	21	65	142,268	142,259	0.06	142,259	22
	8/8/8	140,466.9	140,599.3	29	48	140,691	(141,114)	0.07	140,658	29
T20_20	7/7/7	1,288,847	1,291,178	14:07	698	-	1,294,782	0.36	1,291,008	14:09
	8/8/8	1,069,069.1	1,069,071	4:05	380	-	1,091,571	2.10	1,070,088	4:11
	10/8/8	689,383.2	689,383.2	3:55	331	-	696,917	1.09	690,868	4:02
	11/8/8	555,465.8	555,465.8	4:00	313	564,298	563,124	1.38	555,959	4:03
	12/8/8	518,052.4	518,064	2:04	169	527,071	525,383	1.41	518,285	2:06
	9/9/9	871,824.2	872,009	4:16	319	-	893,495	2.46	872,521	4:20
	10/10/10	689,032.6	689,032.6	4:24	252	692,042	(702,191)	0.44	690,826	4:28
	11/11/11	555,083.7	555,083.7	4:42	219	564,608	561,117	1.09	555,115	4:43
	12/12/12	517,556.9	517,691.5	3:21	131	517,942	(523,999)	0.05	517,815	3:22

Problem	Capacity	Lagrange LB	LP Master	Time LP	Number Columns	Lagrange Heuristic	Fixing Heuristic	Gap (%)	IP	Time IP
T20_30	10/7/7	904,564.2	905,093.8	17:03	768	-	915,156	1.13	908,778	18:31
	11/7/7	688,115	688,305	15:42	727	-	793,160	15.24	793,160	3:00:38*
	8/8/8	1,794,300.2	1,794,475.2	6:01	445	-	1,799,225	0.27	1,796,266	6:22
	9/7/6	1,373,846.8	1,374,543.6	29:04	966	-	1,429,496	4.02	1,419,311	3:40:20*
	9/9/9	1,260,269.8	1,260,326.4	6:26	350	-	1,275,671	1.22	1,270,337	11:39
	10/10/10	887,837	888,082	5:19	250	888,082	888,082	0.01	888,082	5:19
	11/11/11	671,287.4	671,297.9	5:30	245	-	735,835	9.61	672,772	5:33
	12/12/12	627,448.7	628,348.2	4:34	113	628,623	(647,015)	0.05	628,447	4:35
T30_10	5/5/5	253,059.3	253,170.2	2:09	260	-	260,536	2.92	257,719	2:15
	6/5/5	216,830.3	216,958.9	1:47	275	-	231,140	6.55	221,516	1:49
	7/5/5	191,038.8	191,139.7	2:07	226	239,412	196,040	2.57	195,640	3:54
	8/5/5	182,608.3	182,638.7	1:33	186	226,306	187,539	2.69	187,129	1:38
	9/5/5	181,545.5	181,630.7	1:27	171	222,501	186,531	2.71	186,121	1:34
	6/6/6	182,487.1	182,487.3	28	189	-	191,768	5.09	183,462	28
	7/7/7	153,001.5	153,001.5	25	120	153,171	153,089	0.06	153,089	25
	9/6/6	147,086.5	147,159	31	92	147,159	147,159	0.00	147,159	32
	9/6/7	145,304	145,304	17	74	145,304	145,304	0.00	145,304	18
	9/7/7	143,442.8	143,549.8	24	73	143,552	(143,598)	0.04	143,552	24
T30_20	8/5/5	2,387,769.5	2,388,310.2	27:39	873	-	2,413,063	1.05	2,400,933	1:20:54
	10/10/10	1,119,272.9	1,119,315.8	3:48	353	1,154,364	1,127,514	0.73	1,122,206	3:55
	14/8/8	569,004.6	571,346.9	15:05	764	-	586,499	2.84	586,499	2:45:33*
	15/8/8	535,599.7	538,203.2	12:40	668	-	542,107	0.86	541,277	2:43:12*
	16/8/8	533,191.6	533,308.6	7:01	443	-	556,767	4.41	536,160	2:15:44
	16/9/9	530,571.7	530,576.7	2:50	167	530,612	(534,659)	0.01	530,612	2:52
	17/10/10	529,709.6	529,709.6	2:23	93	530,078	530,027	0.06	529,745	2:24
T30_30	14/8/8	708,663.5	709,580	18:42	864	-	772,204	8.90	772,204	3:42:47*
	15/8/8	656,288.5	656,588.5	15:32	694	-	692,793	5.52	692,793	3:22:30
	16/8/8	649,225.3	649,369.2	8:14	487	-	668,107	2.89	668,107	2:58:16*
	16/9/8	647,330.6	647,333.1	5:10	288	-	(657,540)	1.48	648,198	9:51
	16/9/9	646,375.9	646,377.8	2:59	217	-	(655,572)	0.63	646,470	3:02

*: pivot limit of 9,999,999 exceeded

() : LP rounding heuristic is worse than either the initial heuristic or the Lagrange heuristic

- : no better integer solution was found during the subgradient optimization phase.

TABLE 7 Comparison of average computation times and number of columns added for the three algorithms

Problem	Average LP time			Average Number of Columns		
	M1	M2	M3	M1	M2	M3
T10_10	7	5	3	98.5	80.7	47.3
		29.2%	47.7%		18.1%	52.0%
T10_20	36:38	10:16	4:16	980.6	554.2	199.2
		72.0%	88.3%		43.5%	79.7%
T10_30	1:29:39	38:15	14:57	1,565.2	1,039.	414.2
		57.4%	83.3%		33.6%7	73.5%
T20_10	1:38	1:13	46	163.8	146.0	110.5
		25.3%	52.7%		10.9%	32.5%
T20_20	7:58	5:50	4:60	506	441.3	312.4
		26.8%	37.3%		12.8%	38.3%
T20_30	24:26	16:07	11:13	871.6	745.0	483.0
		34.1%	54.1%		14.5%	44.6%
T30_10	3:04	2:00	1:07	264.8	232.3	166.6
		34.6%	63.5%		12.3%	37.1%
T30_20	18:39	11:51	10:13	713.4	620.0	480.1
		36.7%	45.2%		13.1%	32.7%
T30_30	19:25	13:30	10:08	732	658.6	510.0
		30.4%	47.8%		10.0%	30.3%

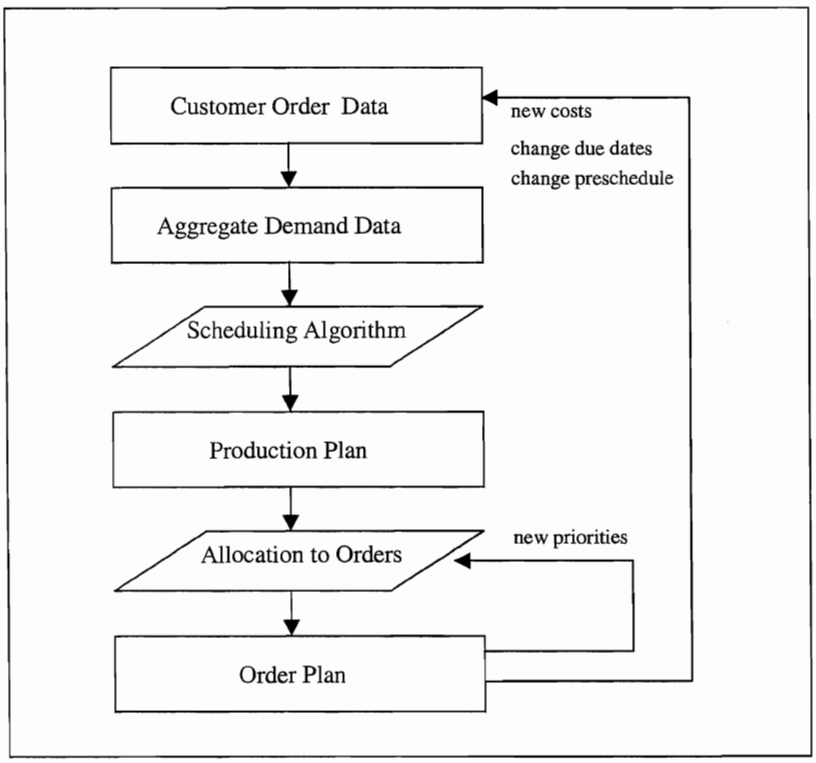
TABLE 8 Comparison of heuristic and IP solutions for the three algorithms

Problem	Average Gap Best Heuristic Solution (%)			Average Gap IP Solution (%)		
	M1	M2	M3	M1	M2	M3
T10_10	0	0.05	0	0	0.05	0
T10_20	3.18	0.32	0.23	0.20	0.31	0.21
T10_30	1.28	0.66	0.38	0.13	0.55	0.26
T20_10	1.83	1.64	1.21	0.89	0.94	0.64
T20_20	1.22	0.99	1.15	0.08	0.28	0.10
T20_30	4.50	4.38	3.95	0.68	3.45	2.51
T30_10	3.12	1.63	2.26	1.40	0.74	1.19
T30_20	1.37	1.11	1.42	0.86	0.71	0.70
T30_30	3.40	3.64	3.88	3.01	1.81	3.49
Average	**2.21**	**1.60**	**1.61**	**0.81**	**0.98**	**1.01**

6. Model impact

Management is interested in quickly generating various alternative high quality schedules from which they can choose. Consequently the speed issue is an important factor. This means that the company does not want an optimal solution at the cost of a large CPU time. Our search for heuristic procedures is hence motivated by practical considerations. The column generation algorithm combined with the Lagrange relaxation for solving the master seems to fit best these conditions. If we stop after the LP fixing heuristic, which takes typically less than half a second after solving the relaxation, we get the fastest results with a small gap. Anyhow, the production schedules thus generated could not be improved manually by people with several years of planning experience. Even for modest sized initial test problems, the planners were unable to find better production plans even after several hours of experimentation. The managers agreed immediately that this was indeed the best solution.

Figure 2 The Production Planning Process



The scheduling algorithm is at the core of the order and production planning process, as can be seen from Figure 2. Starting from the customer order data, the demand data is aggregated according to tire type and due date. First we calculate the net demand for every s.k.u. by subtracting the available inventory from the original demand. At this point we also take the production from the dedicated lines into account. We will hence have no beginning inventory and the size of the demand data is reduced.

Subsequently, the remaining demand data is aggregated over all the tires which make use of the same molds. The scheduling algorithm generates a production plan. The last step is to allocate the production to orders, based on priorities. Within each pool of orders with the same due date, different priorities are allocated to different orders. This gives an order plan with a possible shipping date and the backorder situation for each order. We will clarify this order allocation with a small example based on the production plan in Table 2 and the demand in Table 3. From Table 3 it is clear that we only have backlog in periods 4 and 6. Consequently we will only look at the orders for periods 4 to 6. Suppose that we have customer order data for these days as given in Table 9. Each order has an order number, a due date and a priority. Within the pool of orders with the same due date, the order with priority 1 is the most important and will be completed first. The next important order has priority 2 etc. Each order consists of a specific demand for one or more tires. If these orders are aggregated per due date and tire type, we get exactly the aggregated demand of Table 3. Allocating the production, as given in Table 3, to each order, taking into account the priorities, will lead to an order plan. The allocation itself works as follows. First, the earliest shipping date is calculated for the item with the highest priority. This earliest shipping date for the associated order is the maximum of the earliest shipping dates for each item in that order. Given this shipping date, the demand for the entire order is filled with the latest possible production. When we look at order 214, for example, we see that it has the highest priority of the orders with due date 6. We first determine the earliest shipping date. From the production plan of Table 3, we observe that for order 214 there will be a backlog of tire 3. Consequently, the earliest shipping date will therefore be day 7. This implies also that the other tires of order 214 only have to be finished by day 7. Hence we allocate the inventory to order 214 as late as possible. The inventory is then updated accordingly and we continue with the order with the next highest priority. This is order 208. By allocating the inventory to order 214 as late as possible, given the shipping date, we eliminate the backlog problem of tire 4 for order 208. If the planners want to split an order, i.e. items within one order can have a different shipping date, than this procedure must be adapted. The order plan is given in Table 10. For each order it is indicated how the production is allocated to that order. We see that orders 136 and 214 cannot be shipped on their due date.

Table 9 Customer order data for small example

Due Date	4		5		6	
Order Nr.	128	136	188	190	208	214
Priority	1	2	1	2	2	1
Tire 1	-	-	100	70	-	-
Tire 2	40	33	-	-	20	5
Tire 3	-	-	16	-	-	39
Tire 4	-	-	6	60	40	12
Tire 5	-	-	-	18	-	-

Table 10 Order plan for small example

Order	128	136		188	190	208	214	
Period	4	4	5	5	5	6	6	7
Tire 1	-	-	-	100	70	-	-	-
Tire 2	40	30	3	-	-	20	5	-
Tire 3	-	-	-	16	-	-	37	2
Tire 4	-	-	-	6	60	40	3	9
Tire 5	-	-	-	-	18	-	-	-

Another concern was the control of important parameters by the planners. The four important parameters here are the set up and backlog cost, the due dates, the preschedules and the order priorities. From Figure 2 we see that they have an impact at different levels of the planning process. Starting from the same production plan, different order plans can be generated by varying the priorities of the orders. By varying the ratio of the set up cost to the backlog cost, the planners can generate various alternative production schedules. This requires a rerun of the scheduling algorithm. A planner can also start from some different preschedules. Finally, the impact of putting in new orders or changing the due dates of old orders can be investigated. This also requires a rerun of the scheduling algorithm.

The algorithm can be used on a daily basis in a rolling horizon scheme. Parallel computers can be used to compute schedules for different tire groups. The frequently updated schedules are also a great aid for better due date promising for the order processing department.

Conclusions

In this paper we have described a column generation based heuristic for an industrial capacitated lot sizing problem. The problem is complicated with set up times and multiple non-identical alternative machines. The algorithm consists of four major parts: 1. Initial upper bound calculation; 2. Calculation of the LP relaxation with column generation; 3. An LP fixing heuristic; 4. Branch-and-bound over the columns generated. We have solved a number of real life problems, varying in size and capacity tightness.

In practice, management is interested in finding various good solutions in a short computation time. This suggests the use of the column generation scheme, together with the Lagrange relaxation for solving the master. This gives the fastest results. Even if we stop after the LP fixing heuristic, we still obtain solutions with a small gap. As such, the algorithm will be a great aid both in production planning and order promising.

Future research will focus on implementing a branch-and-price algorithm in order to find the true IP optimum. In one possible scheme we could branch on the original variables x_{ikt}. It is possible to adapt the DP algorithm for the subproblem so that it does not generate columns which violate previous branching constraints by setting limits on the production state variables m. It might be useful to look at other heuristics to find good solutions. Finally, it is worthwhile to investigate the impact of the additional issues as discussed in Section 4.

Appendix

Derivation of the Lagrange objective function:
Dualizing constraint (10) gives:

$$\sum_{i\in T}\sum_{q\in Q_i}\cos t_{iq}\, z_{iq} - \sum_{k\in H}\sum_{t\in P}\pi_{kt}\,(nh_{kt} - \sum_{i|(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\, z_{iq})$$

$$= \sum_{i\in T}\sum_{q\in Q_i}\cos t_{iq}\, z_{iq} + \sum_{t\in P}\sum_{k\in H}\sum_{i|(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\, z_{iq} - \sum_{k\in H}\sum_{t\in P}\pi_{kt}\, nh_{kt}$$

$$= \sum_{i\in T}\sum_{q\in Q_i}\cos t_{iq}\, z_{iq} + \sum_{t\in P}\sum_{(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\, z_{iq} - \sum_{k\in H}\sum_{t\in P}\pi_{kt}\, nh_{kt}$$

$$= \sum_{i\in T}\sum_{q\in Q_i}\cos t_{iq}\, z_{iq} + \sum_{t\in P}\sum_{i\in T}\sum_{k|(i,k)\in TXH}\sum_{q\in Q_i} a_{iktq}\, z_{iq} - \sum_{k\in H}\sum_{t\in P}\pi_{kt}\, nh_{kt}$$

$$= \sum_{i\in T}\sum_{q\in Q_i}(\cos t_{iq} + \sum_{t\in P}\sum_{k|(i,k)\in TXH} a_{iktq})\, z_{iq} - \sum_{k\in H}\sum_{t\in P}\pi_{kt}\, nh_{kt}$$

Acknowledgements

We would like to thank Mr. Pierre Pringiers, chairman of Solideal and Mr. Manu Jans for bringing this interesting problem to our attention and for the good and fruitful cooperation.

References

BARNHART, C., E.L. JOHNSON, G.L. NEMHAUSER, M.W.P. SAVELSBERGH, AND P.H. VANCE. 1998. Branch-and-Price: Column generation for solving huge integer programs. *Operations Research*. 46 (3), 316-329.

BELVAUX, G., AND L.A. WOLSEY. 1998. Lot-sizing problems: modelling issues and a specialized branch-and-cut system BC-PROD. *CORE discussion paper*, 9849, 34 p.

CATTRYSSE, D., M. SALOMON, R. KUIK, AND L.N. VAN WASSENHOVE. 1993. A dual ascent and column generation heuristic for the discrete lotsizing and scheduling problem with setup times. *Management Science*. 39 (4), 477-486.

DEGRAEVE, Z., AND L. SCHRAGE. 1997. A tire production scheduling system for Bridgestone/ Firestone Off-The-Road. *Operations Research*. 45 (6), 789-796.

DEGRAEVE, Z., AND L. SCHRAGE. 1998. HOP: A software tool for production scheduling at Bridgestone/Firestone Off-The-Road. *European Journal of Operational Research*. 110 (2), 188-198.

DIABY, M., H.C. BAHL, M.H. KARWAN, AND S. ZIONTS. 1992. A Lagrangean relaxation approach for very-large-scale capacitated lot-sizing. *Management Science*. 38 (9), 1329-1340.

EPPEN, G.D., AND R.K. MARTIN. 1987. Solving multi-item capacitated lot-sizing problems using variable redefinition. *Operations Research*. 35 (6), 832-848.

FISHER, M.L. 1985. An applications oriented guide to Lagrangian relaxation. *Interfaces*. 15 (2), 10-21.

KANG, S., K. MALIK, AND L.J. THOMAS. 1999. Lotsizing and Scheduling on Parellel Machines with Sequence-Dependent Setup Costs. *Management Science*. 45 (2), 273-289.

LASDON, L.S., AND R.C. TERJUNG. 1971. An efficient algorithm for multi-item scheduling. *Operations Research.* 19, (4), 946-969.

MAGNANTI, T.L., AND R. VACHANI. 1990. A strong cutting plane algorithm for production scheduling with changeover costs. *Operations Research.* 38 (3), 456-473.

MANNE, A.S. 1958. Programming of economic lot sizes. *Management Science.* 4 (2), 115- 135.

MARTIN, R.K. 1999. Large Scale Linear and Integer Optimization: A Unified Approach. Kluwer Academic Publishers, Boston, p. 740.

SCHRAGE, L. 1982. The multiproduct lot scheduling problem. In *Deterministic and Stochastic Scheduling*, M.A.H. Dempster, J.K. Lenstra and A.H.G. Rinnooy Kan (eds.). NATO Advanced Study Institutes Series, D. Reidel, Holland, 233 – 244.

SCHRAGE, L. 1995. LINDO: Optimization software for Linear Programming. Lindo systems Inc., Chicago, IL.

TRIGEIRO, W., L.J. THOMAS, AND J.O. MCCLAIN. 1989. Capacitated lot sizing with set-up times. *Management Science.* 35 (3), 353-366.

VANDERBECK, F. 1998. Lot-sizing with start-up times. *Management Science.* 44 (10), 1409-1425.

WOLSEY, L.A. 1998. Integer programming. Wiley-Interscience series in discrete mathematics and optimization. p. 264.