Omega 39 (2011) 627–633

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

A column generation heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint

Horst Tempelmeier *

Department of SCM and Production, Albertus Magnus-Platz, D-50932 Köln, Germany

ARTICLE INFO

Article history:
Received 5 June 2010
Accepted 14 January 2011
Processed by B. Lev
Available online 1 February 2011

Keywords:
Production planning and control
Heuristics
Inventory control

ABSTRACT

This paper deals with the dynamic multi-item capacitated lot-sizing problem under random period demands (SCLSP). Unfilled demands are backordered and a fill rate constraint is in effect. It is assumed that, according to the static-uncertainty strategy of Bookbinder and Tan [1], all decisions concerning the time and the production quantities are made in advance for the entire planning horizon regardless of the realization of the demands. The problem is approximated with the set partitioning model and a heuristic solution procedure that combines column generation and the recently developed ABC_β heuristic is proposed.

1. Introduction

We consider the stochastic version of the dynamic multi-item capacitated lot-sizing problem (CLSP). The problem is to determine production quantities to satisfy demands for multiple products over a finite discrete time horizon such that the sum of setup and holding costs is minimized, whereby a capacity constraint of a resource must be taken into consideration. In contrast to the deterministic CLSP, we assume that for every product k and period t the demand is a random variable D_{kt} ($k=1,2,\ldots,K$; $t=1,2,\ldots,T$). The period demands are non-stationary (to permit dynamic effects such as seasonal variations, promotions, or general mixtures of known customer orders with random portions of period demands), which usually is the case in a material requirements planning (MRP) based environment. Demand that cannot be filled immediately from stock on hand is backordered. As the precise quantification of shortage penalty costs which involve intangible factors such as loss of customer goodwill is very difficult, if not impossible, we assume that management has specified a target service level. In particular, we assume that the fill rate criterion (β service level) is in effect, as this criterion is very popular in industrial practice (see Tempelmeier [2]).

Industrial (MRP based) planning practice usually applies a forecasting procedure that provides a deterministic time series of expected future demands. Uncertainty is taken into consideration by reserving a fixed amount of inventory as safety stock (see Wortmann [3], Baker [4]). The amount of this reserve stock is usually computed with simple rules of thumb borrowed from stationary inventory theory, e.g. the standard deviation of the demand during the risk period is multiplied by a quantile of the standard normal distribution. In this way, it is almost impossible to meet targeted service levels. In addition, using time-independent safety stocks under dynamic conditions may result in significant cost penalties (see Tunc et al. [5]).

It is obvious that apart from the MRP-inherent neglect of limited capacities this widely used approach completely ignores the impact that lot sizes have on the absorption of risk. For example, in a case when due to high setup costs large lot sizes are used which cover the demands of many periods, it probably will be optimal to use no safety stock at all. On the other hand, if setup times or costs are reduced through technical measures in order to reduce lot sizes and the associated cycle stock, the required safety stock will increase.

In addition, which is even more problematic, the dynamic alteration of the materials requirements as a consequence of newly observed demand realizations according to the MRP planning process leads to random releases of production lots, as the actual timing and size of the required replenishments are the outcome of the demand process, which is random.

The resulting increase of the variance of the production quantities may have some unwanted consequences. First, in multi-level bill-of-material structures (or supply chains), the random change of a production order of a parent item leads to random requirements for its predecessors. This may cause the rescheduling of the production orders for the predecessors, which is a problem if a predecessor comes from an external supplier. If production orders are rescheduled, then demand variations occur

* Tel.: +49 221 470 3378; fax: +49 221 470 5140.
E-mail address: tempelmeier@wiso.uni-koeln.de

0305-0483/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.omega.2011.01.003

that are propagated upstream through the supply chain, and which must be accounted for through buffers. In the literature this issue is discussed as planning nervousness. Second, the random change of the timing or size of a production lot directly translates into random resource requirements. For a machine, this is usually not a problem as long as the capacity of the machine is not overloaded. If an overload occurs, however, with fixed machine capacities this implies that the production plan becomes infeasible. In this case the planned due dates will be missed. This is one of the biggest problems found in short-term production planning in industry. In addition, there may even be cases when due to technical constraints the production quantities are unchangeable. This is often true in the process industries. Finally, if the considered resource is a human operator, then it may be unfavorable or even prohibited by labor agreement to change the workload in a period.

One countermeasure is the definition of a planning horizon with an unchangeable production plan (frozen schedule). This is what we study in the current paper. In the following, we assume that, according to the static-uncertainty strategy of Bookbinder and Tan [1], all decisions concerning the time and the production quantities are made in advance for the entire planning horizon, which is equivalent to using a frozen schedule. The unavoidable randomness of demand is accounted for through the appropriate sizing of the orders. Other than Bookbinder and Tan [1], we consider multiple products, a resource with limited capacity and a fill rate constraint.

The rest of this paper is organized as follows. In Section 2 the relevant literature is reviewed. Next, in Section 3, the considered stochastic lot-sizing problem under a fill rate per cycle constraint as proposed by Tempelmeier and Herpers [6] is approximated with a set partitioning model. Then, in Section 4, we present a heuristic column generation procedure to solve the LP-relaxation of this model and combine this procedure with the ABC_β heuristic proposed in Tempelmeier and Herpers [6] to solve the complete problem. The results of a numerical experiment are reported in Section 5. Finally, Section 6 contains some concluding remarks.

2. Literature

The deterministic multi-item dynamic capacitated lot-sizing problem has been studied for a long time. For recent overviews see Karimi et al. [7], Jans and Degraeve [8], Robinson et al. [9] and Buschkühl et al. [10]. However, only a limited number of researchers have considered dynamic capacitated lot sizing under random demand. A literature overview is presented in Sox et al. [11]. Sox and Muckstadt [12] solve a variant of the stochastic dynamic CLSP, where item- and period-specific backorder costs as well as extendible production capacities are considered. These authors propose a Lagrangean heuristic to solve the resulting non-linear integer programming problem that is repeatedly applied in a dynamic planning environment. Martel et al. [13] develop a branch-and-bound procedure for the solution of a similar model formulation.

Brandimarte [14] considers the stochastic CLSP where the uncertainty of the demand is represented through a scenario tree. In this case, the period demands are modeled as discrete random variables. The evolution of demand over time is depicted with a directed layered tree, where each layer corresponds to a planning period and the nodes are linked to realizations of the discrete stochastic demand process. The resulting large-scale deterministic MILP model is then solved with a commercially available solver using rolling schedules with lot-sizing windows. As demonstrated by Brandimarte [14], the scenario-based approach suffers from a dramatically increasing complexity, if the number of periods and/or the number of possible outcomes of the period demands are increased. In addition, currently there are no scenario-based models available which could account for product-specific fill rate constraints.

Tempelmeier and Herpers [6] propose a formulation of the dynamic capacitated lot-sizing problem under random demand, when the performance is measured in terms of a fill rate per cycle which is a popular performance measure in industry. They propose the ABC_β heuristic which is an extension of the $A/B/C$ heuristic of Maes and Van Wassenhove [15].

3. Problem formulation

Below the following notation is used:

b_t	capacity in period t (time units)
$\beta_k^\star$	target fill rate per cycle for product k
c_n	total cost of production plan n of product k
D_{kt}	demand for product k in period t
F_{kt}	backorders of product k in period t
δ_n	binary selection variable for plan n
γ_{kt}	binary setup indicator for product k in period t
h_k	inventory holding cost per time period per unit of product k
I_{kt}	net inventory for product k at the end of period t
$I_{kt}^{f,\mathrm{end}}$	backlog of product k at the end of period t
$I_{kt}^{f,\mathrm{prod}}$	backlog of product k after production in period t, but before demand occurrence
K	number of products
κ_{nt}	capacity requirement of production plan n in period t
l_{kt}	number of periods since the last setup (product k, period t)
M	sufficiently large number
ω_{kt}	indicator variable: $\omega_{kt}=1$, if production of product k takes place in period t+1; $\omega_{kt}=0$ otherwise
$\mathcal{P}_k$	set of production plans for product k
π_t	dual variable associated to the capacity requirement constraint of period t
q_{kt}, q_{nt}	lot size of product k (production plan n) in period t
σ_k	dual variable associated to the plan selection constraint for product k
s_k	setup costs for product k
tb_k	capacity usage for production of one unit of product k
T	length of the planning horizon
$[x]^+$	$\max\{0,x\}$
$[x]^-$	$\min\{0,x\}$

Consider K products that are produced to stock on a single resource with given period capacities b_t (t=1,2,...,T). The planning situation is completely identical with that assumed in the classical dynamic capacitated lot-sizing problem (CLSP) without setup times. However, there is one exception: For each product k and each time period t, the period demands D_{kt} are random variables with a known probability distribution and given period-specific expected values $E\{D_{kt}\}$ and variances $V\{D_{kt}\}$. These data, which in a dynamic planning environment vary over time, are the outcome of a forecasting procedure. The demands of the various products are mutually independent and there is no autocorrelation. Unfilled demands are backordered. The amount of backorders is controlled by imposing a fill rate per cycle constraint. We define the fill rate per cycle as the ratio of the expected demand observed during the coverage time of a production order that is routinely filled from available stock on hand and the actual lot size. More precisely, let τ be a production period of product k, let t be the last period before the next production of product k takes place and let $q_{k\tau}$ be the lot size produced in period τ covering the demand up to period t. Finally, let F_{ki} be the backorders of product k that occur during period i. Then for a target service level β_k^* it is required, that

$$1-\frac{E\left\{\sum_{i=\tau}^{t} F_{ki}\right\}}{E\left\{\sum_{i=\tau}^{t} D_{ki}\right\}} \geq \beta_k^* \quad k=1,2,\ldots,K \tag{1}$$

This constraint is equivalent to the fill rate definition under stationary conditions which relates the average backorders per cycle to the average replenishment quantity (see Silver and Bischak [16]). However, it is a sharper requirement, as not only in the long run, but also in each production cycle that the fill rate target must be met. At the beginning of the planning horizon there is a known initial inventory I_{k0} $(k=1,2,\ldots,K)$.

As noted above, our planning approach implements the "static uncertainty" strategy of Bookbinder and Tan [1] which means that at the beginning of the planning horizon, the complete production plan is fixed in advance, including the timing and the size of the production quantities. The resulting dynamic multi-item stochastic capacitated lot-sizing problem is represented by the following model SCLSP$_\beta$ (see Tempelmeier and Herpers [6]):

Model SCLSP$_\beta$

$$\text{Minimiere } Z=\sum_{k=1}^{K}\sum_{t=1}^{T}(s_k\cdot\gamma_{kt}+h_k\cdot E\{[I_{kt}]^+\}) \tag{2}$$

subject to

$$I_{k,t-1}+q_{kt}-D_{kt}=I_{kt} \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{3}$$

$$q_{kt}-M\cdot\gamma_{kt}\leq 0 \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{4}$$

$$\sum_{k=1}^{K} tb_k\cdot q_{kt}\leq b_t \quad t=1,2,\ldots,T \tag{5}$$

$$I_{kt}^{f,\text{prod}}=-[I_{k,t-1}+q_{kt}]^- \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{6}$$

$$I_{kt}^{f,\text{end}}=-[I_{kt}]^- \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{7}$$

$$F_{kt}=I_{kt}^{f,\text{end}}-I_{kt}^{f,\text{prod}} \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{8}$$

$$l_{kt}=(l_{k,t-1}+1)\cdot(1-\gamma_{kt}) \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{9}$$

$$l_{k,0}=-1 \quad k=1,2,\ldots,K \tag{10}$$

$$\omega_{kt}=\gamma_{k,t+1} \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T-1 \tag{11}$$

$$\omega_{kT}=1 \quad k=1,2,\ldots,K \tag{12}$$

$$1-\frac{E\left\{\sum_{j=t-l_{kt}}^{t} F_{kj}\right\}}{E\left\{\sum_{j=t-l_{kt}}^{t} D_{kj}\right\}} \geq \beta_k^* \quad k=1,2,\ldots,K;\ t\in\{t|\omega_{kt}=1\} \tag{13}$$

$$q_{kt}\geq 0 \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{14}$$

$$\gamma_{kt}\in\{0,1\} \quad k=1,2,\ldots,K;\ t=1,2,\ldots,T \tag{15}$$

The objective function (2) minimizes the total setup costs and expected inventory holding costs, where $[I_{kt}]^+$ is the inventory on hand at the end of period t for product k. Eq. (3) is the standard inventory balance equation. Constraint (4) forces the setup indicator γ_{kt} to 1, whenever there is a positive production quantity q_{kt}, according to the assumptions of a big-bucket lot-sizing model. Constraint (5) requires that the available capacity b_t per period must not be exceeded. Eq. (6) defines the backlog in period t immediately after a production has taken place and all outstanding backorders, if any, have been filled as much as possible before the new demand of period t occurs. Eq. (7) describes the backlog at the end of period t. Eq. (8) defines the backorders that newly occurred in period t.

The remaining equations are used for book-keeping. In order to calculate the average fill rate during an order cycle, we count the number of periods covered by a production quantity with the help of variable l_{kt}. Eqs. (11) and (12) set the indicator variable ω_{kt} to 1, if either period $t+1$ is a setup period or the planning horizon ends in period t. In addition, the length of an order cycle (the number of periods between two consecutive setups) must be known. This is computed by means of Eqs. (9) and (10). l_{kt} is reset to zero whenever $\gamma_{kt}=1$, i.e., when t is a setup period for product k. Otherwise, l_{kt} is incremented by one to $(l_{k,t-1}+1)$. Eq. (13) defines the expected fill rate within the actual production cycle since the last production of product k.

Under deterministic demand the model reduces to the standard formulation of the CLSP. In this case all random variables are deterministic, and constraints (6)–(13) can be omitted.

As originally proposed for a variant of the deterministic CLSP by Manne [17], the lot-sizing problem can be approximated by a set partitioning model as follows. Define for each product k a set $\mathcal{P}_k$ $(k=1,2,\ldots,K)$ of alternative production plans over the planning horizon T. Each production plan n is composed of given setup periods and associated lot sizes that cover an integer number of period demands under consideration of the fill rate constraint. The computation of the lot sizes is discussed below.

As the production quantities are fixed in advance, with a given production plan n of product k, the expected total setup and holding costs c_n and the exact capacity requirements κ_{nt} in period t $(k=1,2,\ldots,k; n\in P_k; t=1,2,\ldots,T)$ can be determined. The problem is then to select for each product exactly one production plan alternative such that in all periods the capacity constraint is respected. The resulting set partitioning model formulation is

Model SCLSP$_{\text{SPP}}$

$$\text{Minimize } Z=\sum_{k=1}^{K}\sum_{n\in P_k} c_n\cdot\delta_n \tag{16}$$

subject to

$$\sum_{k=1}^{K}\sum_{n\in P_k}\kappa_{nt}\cdot\delta_n\leq b_t \quad t=1,2,\ldots,T\ (\pi_t) \tag{17}$$

$$\sum_{n\in P_k}\delta_n=1 \quad k=1,2,\ldots,K\ (\sigma_k) \tag{18}$$

$$\delta_n\geq 0 \quad k=1,2,\ldots,K; n\in P_k \tag{19}$$

The objective function (16) minimizes the sum of the expected costs of all selected production plans. δ_n is a binary variable that selects production plan $n\in\mathcal{P}_k$. Constraint (17) ensures that the period capacity of the resource in period t is respected, whereby κ_{nt} is the capacity requirement resulting from production plan n in period t. Eq. (18) states that for each product exactly one production plan must be selected.

While model SCLSP$_{\text{SPP}}$ looks exactly like the standard SPP formulation found in the literature (see Lasdon and Terjung [18], Cattrysse et al. [19], Haase [20]), the basic difference is hidden in the cost coefficients c_n. The coefficient c_n which represents the expected costs of production plan $n\in\mathcal{P}_k$ is the result of an embedded optimization problem. This problem consists in determining for each given setup period the minimum lot size required to achieve the target service level at the end of the associated production cycle.

Assume that for production plan $n\in\mathcal{P}_k$ a setup pattern is given with J_n setups in periods $\tau_j (j=1,2,\ldots J_n)$. Then for each setup j the problem is to find the minimum lot size $q_{n\tau_j}$, that respects the service level constraint for the current production cycle j. This problem can be stated as follows:

Model MINQ$_j$

$$\text{Minimize } q_{n\tau_j} \tag{20}$$

s.t.

$$1-\frac{E\left\{\sum_{i=\tau_j}^{\tau_{j+1}-1} F_{ni}(q_{n\tau_j})\right\}}{E\left\{\sum_{i=\tau_j}^{t} D_{ki}\right\}} \geq \beta_k^{\star} \tag{21}$$

with $\tau_{J_n+1}=T+1$. The minimum lot size is calculated with the help of a search procedure. For a given setup pattern with J_n setups and associated lot sizes the expected costs c_n of production plan $n \in \mathcal{P}_k$ can be easily computed. The setup cost is $s_k \cdot J_n$. The resulting expected on hand inventory at the end of period t is derived as follows. Let

$$Q_n^{(t)} = I_{k0} + \sum_{i=1}^{t} q_{ni} \tag{22}$$

be the sum of the initial inventory I_{k0} and the production quantities produced from period 1 up to period t, according to production plan n, and let

$$Y_k^{(t)} = \sum_{i=1}^{t} D_{ki} \tag{23}$$

denote the cumulated demands of product k from period 1 up to period t. Then the expected inventory on hand at the end of period t is equal to $E\{I_{nt}^p\} = E\{[Q_n^{(t)} - Y_k^{(t)}]^+\}$, where $[X]^+ = \max[X,0]$. This can be written as

$$E\{I_{nt}^p\} = Q_n^{(t)} - E\{Y_k^{(t)}\} + G^1_{Y_k^{(t)}}(Q_n^{(t)}) \quad t=1,2,\ldots,T, \tag{24}$$

where $G_Y^1(Q) = \int_Q^\infty (y-Q)\cdot f_Y(y)\cdot dy$ is the first-order-loss function with respect to the random variable Y and the quantity Q.

The backorders that newly occur in period t are the difference between the backlog at the end of that period and the backlog that remained after a replenishment at the beginning of the period, before the occurrence of the period demand, i.e.

$$F_{nt}(q_{n\tau_j}) = -[I_{nt}]^- - (-[I_{n,t-1}+q_{n\tau_j}]^-) \quad t=\tau_j,\tau_j+1,\ldots,T, \tag{25}$$

where $[X]^- = \min[X,0]$. If the inventory position I_{nt} at the end of period t is negative, then the backlog is $-I_{nt}$. This is the first term in (25). Similarly, if the sum of the inventory position at the end of period $(t-1)$ plus the replenishment at the beginning of period t is negative, then the backlog is $-[I_{n,t-1}+q_{n\tau_j}]$. This is the second term in (25). The difference is the backorders that newly occurred in period t. Note that the expected backlog at the end of period t is the expected positive difference between the cumulated demand up to period t, $Y_k^{(t)}$ and the cumulated production quantity, $Q_n^{(t)}$, i.e. $G^1_{Y_k^{(t)}}(Q_n^{(t)})$. Similarly, the expected backlog after the replenished at the beginning of period t is equal to the expected positive difference between the cumulated demand up to period $(t-1)$, $Y_k^{(t-1)}$ and the cumulated production quantity up to period t, $Q_n^{(t)}$, i.e. $G^1_{Y_k^{(t-1)}}(Q_n^{(t)})$.

The expected number of backorders associated to period t that occur if production plan n is applied for the considered product can then be written as

$$E\{F_{nt}(q_{n\tau_j})\} = G^1_{Y_k^{(t)}}(Q_n^{(t)}) - G^1_{Y_k^{(t-1)}}(Q_n^{(t)}) \quad t=\tau_j,\tau_j+1,\ldots,T \tag{26}$$

For example, assume a production plan (the index n is now omitted) for T=6 periods with identical period demands that are normally distributed with mean $\mu_t = 100$ and standard deviation $\sigma_t = 30$ (t=1,2,…,6). If two production lots q_1=312.68 and q_4=322.56 are scheduled in periods $\tau_1 = 1$ and $\tau_2 = 4$, then the development of the expected backorders and the expected inventory on hand is shown in Table 1.

Table 1
Production plan.

Period	Lot size	E {Inventory on hand}	E {Backorders}	β
1	312.68	212.68	0.00	–
2	–	112.74	0.05	–
3	–	27.69	14.95	0.95
4	322.56	235.24	0.00	–
5	–	135.79	0.54	–
6	–	50.25	14.46	0.95

For example, the expected backorders in period 3 (14.95) are calculated as follows:

$$\mu_{Y(2)} = 200, \quad \sigma_{Y(2)} = 42.43, \quad \mu_{Y(3)} = 300, \quad \sigma_{Y(3)} = 51.96$$

$$G^1_{Y(2)}(312.68) = 0.05, \quad G^1_{Y(3)}(312.68) = 15.00$$

$$E\{F_3(312.68)\} = 15.00 - 0.05 = 14.95$$

The lot sizes have been selected such that the target fill rate of 95% at the end of each production cycle is met. Although both production cycles have identical lengths, the lot sizes are different, as the net inventory at the end of the first cycle has an impact on the production requirements in the second cycle.

It is known that the above set partitioning formulation is a good approximation of the capacitated lot-sizing problem if the number of products is significantly greater than the number of periods. See Karimi et al. [7]. This is very likely the case in the planning environment that the considered lot-sizing problem is usually applied in. It should be noted that, if appropriately applied in an MRP-based planning environment, the SCLSP is an operational short-term planning problem with only a small number of periods. In addition, the SCLSP is a big-bucket lot-sizing model with a period length of, say, one week and many products produced during the planning horizon. Moreover, the number of products will be significantly greater than the number of planning periods, if there are many products with sporadic demands. Hence, the number of products is usually much larger than the number of periods. For example, in a typical practical planning environment that we observed in industry the planning horizon was six weeks (periods) and the number of products was about 70. If by contrast the planning periods are short, then other model formulations using small buckets are probably more adequate.

4. Solution approach

4.1. Overall procedure

With respect to the number of production plans considered we follow the same approach as the set partitioning based solution procedures proposed for the deterministic CLSP, i.e. we consider only production plans with integer numbers of demand periods covered by any production lot. As stated, for example, by Cattrysse [19], this is a simplifying assumption, as with capacity constraints the optimum solution may comprise a production lot that covers less than the full demand of a period.

It is well known that even with this simplifying assumption the number of production plans becomes prohibitively large even for small problem instances. Therefore, generating all production plans or rather setup patterns in advance is only feasible for very small problem instances. However, as Model SCLSP$_{\text{SPP}}$ has the same formal structure as its deterministic counterpart, we propose to use a column generation approach that generates the

Table 2
Column generation procedure.

(a) Solve the LP-relaxation of the restricted master problem.
Let σ_k $(k=1,2,\ldots,K)$ and π_t $(t=1,2,\ldots,T)$ be the optimal shadows prices.
(b) For $k=1,2,\ldots,K$ (Subproblem k)
Solve the stochastic uncapacitated lot-sizing problem for product k.
Let $\overline{c}_k^{\text{opt}}$ be the minimum objective value.
If $\overline{c}_k=\overline{c}_k^{\text{opt}}-\sigma_k<0$ then
Add a column for the optimal production plan of product k
to the restricted master problem.
endif
endfor
If at least one new column has been added, goto Step (a); otherwise goto Step (c).
(c) Fix the production plans for all products with integer values of the δ−variables and adjust the period capacities accordingly.
Solve the remaining problem with the ABC_β heuristic.

candidate production plans as required. Column generation (see Desaulniers et al. [21], Pochet and Wolsey [22]) is a general iterative solution technique for large-scale linear programs. A column generation procedure starts with a restricted master problem that contains only a few variables. New columns (variables) are generated in an iterative procedure as needed. In each iteration, basically two steps are performed. First, the restricted master problem is solved which provides optimum shadow prices. Second, in order to find the most promising new variable to be introduced into the restricted master problem, a subproblem is solved with the objective to minimize the reduced costs. If the minimum reduced costs are ≥ 0, then there is no improving variable and the original problem is solved.

In the current problem the LP-relaxation of the set partitioning problem serves as the master problem. The corresponding subproblem comprises K product-specific uncapacitated dynamic lot-sizing problems with random demand and a fill rate per cycle constraint. These are solved with the exact solution procedure of Tempelmeier and Herpers [23]. To find an all-integer solution of model $\text{SCLSP}_{\text{SPP}}$ the ABC_β heuristic is applied to a reduced model including all products with fractional δ−variables. See Cattrysse et al. [19], Haase [20]. The overall procedure is specified in Table 2.

As proposed by Haase [20], we start the column generation procedure with a dummy production plan with zero production and prohibitive high costs for each product.

4.2. Solution of a subproblem

A subproblem, i.e. a stochastic uncapacitated lot-sizing problem for product k, can be cast as a shortest-path problem with $T+1$ nodes labeled $(1,2,\ldots,T+1)$, as depicted in Fig. 1 for the case of $T=3$. An edge originating at node τ and ending at node j specifies that the inventory on hand after production in period τ covers the demands from period τ to $j-1$ to the extent dictated by the target fill rate. The next setup is then scheduled for period j.

According to the general structure of a column generation procedure that uses an LP relaxation of model $\text{SCLSP}_{\text{SPP}}$, the costs associated with an edge starting at node τ and ending at node j are given as

$$c_{\tau j}=E\{C_{\tau j}(P_\tau^{\text{opt}})\}-\pi_\tau\cdot tb_k\cdot q_{\tau j}, \tag{27}$$

where tb_k denotes the capacity requirements for one unit of product k. The term $E\{C_{\tau j}(\cdot)\}$ represents the corresponding expected setup and holding costs that occur when the production (setup) in period τ covers the demand up to period $(j-1)$. As the production quantity required in period τ to guarantee the target service level depends on the net inventory at the beginning of

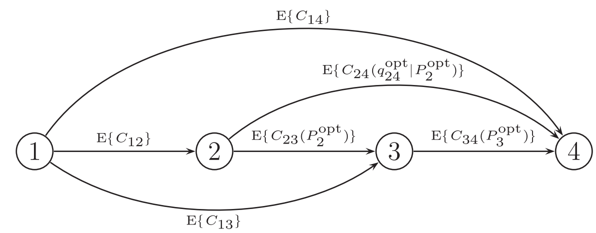

Fig. 1. Shortest-path network.

period τ, which in turn is influenced by the optimum production plan up to that period (denoted as P_τ^{opt}), the costs cannot be computed in advance (as in the deterministic counterpart of the problem) but must be calculated during the solution procedure. The details of the dynamic programming solution procedure for this problem are given in Tempelmeier and Herpers [23].

Let $\overline{c}_k^{\text{opt}}$ denote the objective value of the optimum solution of the shortest-path problem for product k. Then the reduced cost of this optimum production plan for product k are

$$\overline{c}_k=\overline{c}_k^{\text{opt}}-\sigma_k. \tag{28}$$

If $\overline{c}_k<0$, the current production plan for product k is added to the set partitioning model. Once for each product a subproblem has been solved, the next instance of model $\text{SCLSP}_{\text{SPP}}$ is generated and its LP-relaxation is solved. The optimum solution provides new values of the dual variables π_t and σ_k which are then used to generate new product-specific subproblems. The procedure ends when no new production plans are generated.

At this point, all production plans with integer δ_n−variables are fixed and their capacity requirements are subtracted from the available period capacities. For the remaining products with fractional δ_n−variables and the residual period capacities, the heuristic ABC_β procedure proposed by Tempelmeier and Herpers [6] is applied.

5. Numerical results

In order to test the quality of the proposed heuristic we conducted a numerical experiment with two different data sets. In a first step, we generated problem instances with up to 20 periods and up to 70 products. The period demands were assumed to be normally distributed. The parameters of the data set are shown in Table 3.

For each combination of the parameters T, K, $\beta_k^\star$ (the same fill rate is used for all products) and ‘Capacity’ ten replications were generated, whereby the expected demands per product and period were drawn from a continuous uniform distribution. For each product, the coefficient of variation (used for the complete demand series of this product) was selected from a discrete uniform distribution with the possible outcomes {0.15, 0.2, 0.25, 0.3, 0.35}. The TBO-values of the products were randomly drawn from the discrete uniform distribution with the possible outcomes {1,5,10}. The holding costs of product k were calculated as $h_k=2\cdot s/\overline{d_k}\cdot \text{TBO}_k^2$ where $\overline{d_k}$ is the average demand per period and the setup costs were $s=500$ for all products. The period capacity of the resource was calculated as follows. First, the average workload $w=\sum_{k=1}^K\overline{d}_k$ was computed. According to the considered workload scenario {low, medium, high}, the capacity b_t was then set as $b_t^{\text{low}}=1.10\cdot w$, $b_t^{\text{medium}}=1.50\cdot w$, and $b_t^{\text{high}}=2\cdot w$ $(t=1,2,\ldots,T)$. Depending on the target fill rate, the resulting utilizations of the solved problem instances spanned between 47% and 97%. In addition, some combinations of fill rate and capacity resulted in problem instances with capacity over-utilization for which consequently no solution was found. In total, 3240 individual problem instances were generated.

Table 3
Parameters used for the first data set.

Number of periods, T	{10, 20}
Number of products, K	{10, 40, 70}
Fill rate, β_c^*	{0.5, 0.6, …, 0.9, 0.98}
Capacity	{low, medium, high}
Mean period demand	Continuous uniform $U(0, 100)$
Coefficient of variation	Discrete uniform $U(0.15, 0.2, 0.25, 0.3, 0.35)$
Time-between-orders, TBO	Discrete uniform $U(1,5,10)$

Table 4
Results for the first experiment.

	Capacity	Average cost increase of ABC_β		Total (%)
		$T=10$ (%)	$T=20$ (%)	
$K=10$	$1.10 \cdot w$	15.91	14.96	15.44
	$1.50 \cdot w$	12.53	9.30	10.91
	$2.00 \cdot w$	10.77	9.05	9.91
$K=40$	$1.10 \cdot w$	45.74	39.71	42.73
	$1.50 \cdot w$	29.00	23.56	26.28
	$2.00 \cdot w$	19.58	16.28	17.93
$K=70$	$1.10 \cdot w$	51.78	51.72	51.75
	$1.50 \cdot w$	33.20	29.34	31.27
	$2.00 \cdot w$	21.47	18.65	20.06
				25.14

Each problem instance was solved with the proposed column generation heuristic (combined with the ABC_β heuristic for the remaining problem, as described above) and with the ABC_β heuristic of Tempelmeier and Herpers [6] alone. For the latter, we used the parameter combination SH/SM/E which means that the products were sorted according to the setup costs over the holding costs ratio, the cost criterion was the Silver–Meal criterion and the east direction was used.

For 2804 problem instances a feasible solution could be found with both heuristics. The SABC heuristic solved 148 problem instances which could not be solved with the column generation heuristic. The reason may be that the candidate plans considered by the column generation heuristic only cover integer numbers of demand periods. Particularly with low capacity, it may be required to consider demand splitting in order to find a feasible solution.

For 97.95% of all solved problem instances the proposed column generation heuristic found the best solution. For these problems, the solution quality of the column generation heuristic was on the average 25.14% better than with the ABC_β heuristic. Table 4 shows the relative cost increase of the ABC_β heuristic compared to the solution found with the column generation (CG) heuristic, i.e. the ratio $(ABC_\beta/CG-1)$.

On the other hand, in 2.05% of all problem instances (exclusively small problems with 10 products) the ABC_β heuristic performed better. However, in these cases the average cost difference was only 0.09%.

In the second part of the experiment, we used the same parameters as before, with one exception. Before generating the period demand data based on the normal distribution as described above, we used a Bernoulli random variable to decide if a positive demand occurs in that period at all. In particular, for each period and product we proceeded as follows: (1) Generate as continuous uniform $U(0,1)$ random number r. (2) If $r \le p$, generate normal distributed demand in the same way as in the first experiment. In this way, intermittent (sporadic) demand was generated. We used $p=0.3$.

Here, 2682 problem instances were solved with both heuristics, whereby ABC_β solved only 1866 instances. The relative performance of both heuristics is as observed with the first experiment. For 98.07% of all solved problem instances the proposed column generation heuristic found the best solution. For these problems, the solution quality of the column generation heuristic was on the average 13.95% better. Only in 1.93% of all instances the ABC_β heuristic performed better with an average cost difference of 0.07%. The detailed results are shown in Table 5, whereby the empty cells denote the problem classes where the ABC_β heuristic could not find a feasible solution and hence no benchmark solutions are available for the column generation heuristic.

Table 5
Results for second experiment (intermittent demand).

	Capacity	Average cost increase of ABC_β		Total (%)
		$T=10$ (%)	$T=20$ (%)	
$K=10$	$1.10 \cdot w$	25.15	–	25.15
	$1.50 \cdot w$	12.03	29.57	20.80
	$2.00 \cdot w$	9.92	14.97	12.45
$K=40$	$1.10 \cdot w$	–	–	–
	$1.50 \cdot w$	11.38	11.46	11.42
	$2.00 \cdot w$	11.04	10.67	10.86
$K=70$	$1.10 \cdot w$	–	–	–
	$1.50 \cdot w$	11.32	9.95	10.64
	$2.00 \cdot w$	12.08	11.79	11.94
				13.95

It appears that the proposed column generation heuristic performs quite well compared to the ABC_β heuristic. This is particularly true when the number of products is large compared to the number of periods on a standard PC whereby the LP problems where solved with CPLEX. Moreover, the cost difference is largest with low capacity or rather high utilization of the resource.

The computation times for the column generation heuristics ranged between 0.2 and 5 s (the latter for the problems with 20 periods and 70 products). For the problems with $T=20$ periods the ABC_β heuristic was about twice as fast on the average. For the problems with $T=10$ periods the column generation heuristic ran slightly faster than the ABC_β heuristic.

6. Conclusion

In this paper we introduced an approximate model for the single level capacitated lot-sizing problem with dynamic stochastic demand under a fill rate constraint. We proposed to combine a column generation procedure to solve the LP-relaxation of the model with the ABC_β heuristic of Tempelmeier and Herpers [6] to solve the remaining problem. The quality of the solutions is compared to the results found with the application of the ABC_β heuristic of Tempelmeier and Herpers [6] alone. It was found that the proposed heuristic is fast and that it provides solutions that are on the average superior to the ABC_β heuristic.

In addition, the set partitioning model has the significant advantage that due to the model structure it is able to easily include setup times. However, in this case, the solution requires a heuristic for the remaining problem that can also handle setup times. This will be a subject for further research.

References

[1] Bookbinder J, Tan J-Y. Strategies for the probabilistic lot-sizing problem with service-level constraints. Management Science 1988;34:1096–108.
[2] Tempelmeier H. Inventory management in supply networks — problems models, solutions. Norderstedt: Books on Demand GmbH; 2006.
[3] Wijngaard J, Wortmann J. MRP and inventories. European Journal of Operational Research 1985;20:218–93.

[4] Baker KR. Requirements planning. In: Graves S, Rinnooy Kan A, Zipkin P, editors. Handbooks in operations research and management science. Logistics of production and inventory, vol. 4. Amsterdam: North-Holland; 1993. p. 571–627. [chapter 11].
[5] Tunc H, Kilic O, Tarim SA, Eksioglu B. The cost of using stationary inventory policies when demand is non-stationary. Omega 2011;39:410–5.
[6] Tempelmeier H, Herpers S. ABC_{β} — a heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. International Journal of Production Research 2010;48:5181–93.
[7] Karimi B, Fatemi Ghomi SMT, Wilson JM. The capacitated lot sizing problem: a review of models and algorithms. Omega 2003;31:365–78.
[8] Jans R, Degraeve Z. Modeling industrial lot sizing problems: a review. International Journal of Production Research 2008;46:1619–43.
[9] Robinson P, Narayanan A, Sahin F. Coordinated deterministic dynamic demand lot-sizing problem: a review of models and algorithms. Omega 2009;37:3–15.
[10] Buschkühl L, Sahling F, Helber S, Tempelmeier H. Dynamic capacitated lot-sizing problems: a classification and review of solution approaches. OR Spectrum 2010;32:231–61.
[11] Sox C, Jackson P, Bowman A, Muckstadt J. A review of the stochastic lot scheduling problem. International Journal of Production Economics 1999;62:181–200.
[12] Sox C, Muckstadt J. Optimization-based planning for the stochastic lot-scheduling problem. IIE Transactions 1997;29:349–57.
[13] Martel A, Diaby M, Boctor F. Multiple items procurement under stochastic nonstationary demands. European Journal of Operational Research 1995;87(1):74–92.
[14] Brandimarte P. Multi-item capacitated lot-sizing with demand uncertainty. International Journal of Production Research 2006;44(15):2997–3022.
[15] Maes J, Van Wassenhove L. A simple heuristic for the multi item single level capacitated lot sizing problem. OR Letters 1986;4:265–73.
[16] Silver E, Bischak D. The exact fill rate in a periodic review base stock system under normally distributed demand. Omega 2011;39:346–9.
[17] Manne A. Programming of economic lot sizes. Management Science 1958;4:115–35.
[18] Lasdon L, Terjung R. An efficient algorithm for multi-item scheduling. Operations Research 1971;19:946–69.
[19] Cattrysse D, Maes J, Van Wassenhove L. Set partitioning and column generation heuristics for capacitated dynamic lot sizing. European Journal of Operational Research 1990;46:38–47.
[20] Haase K. Solving large-scale capacitated lot-sizing problems close to optimality. Technical Report, Technische Universität Dresden; 2005.
[21] Desaulniers G, Desrosiers J, Solomon MM. Column generation. Springer-Verlag; 2005.
[22] Pochet Y, Wolsey LA. Production planning by mixed integer programming. New York: Springer; 2006.
[23] Tempelmeier H, Herpers S. Dynamic uncapacitated lot sizing with random demand under a fill rate constraint. Technical Report, Universität zu Köln, Seminar für Supply Chain Management und Produktion; 2010.