THE CONTROL OF AN INTEGRATED PRODUCTION-INVENTORY SYSTEM WITH JOB SHOP ROUTINGS AND STOCHASTIC ARRIVAL AND PROCESSING TIMES

P.L.M Van Nyen[1] * – J.W.M. Bertrand[1] – H.P.G. Van Ooijen[1] – N.J. Vandaele[2]
[1] *Technische Universiteit Eindhoven* - [2] *University of Antwerp*
[*] *Tel.: +31 40 247 4388 E-mail: p.v.nyen@tm.tue.nl*

Abstract:
This paper investigates a multi-product multi-workcenter production-inventory system, characterized by job shop routings and stochastic demand interarrival times, set up times and processing times. The stock points and the production system are controlled integrally by a centralized decision maker. We present a decision-making procedure to determine the control parameters that minimize overall relevant costs while satisfying prespecified customer service levels. The procedure is tested in an extensive simulation study and the results are discussed.

Keywords:
Production-inventory system; queueing network analyser; performance analysis; production control; inventory control.

1. Introduction

This paper addresses the problem of determining optimal inventory and production control decisions for an integrated production-inventory (PI) system in which multiple products are made-to-stock through a functionally oriented shop. As can be seen in Figure 1, customer orders arrive to the stock points. The stock points generate replenishment orders that are, in this integrated PI system, equivalent to production orders. In this paper, all stock points are controlled using (R,S) inventory policies. The production orders are manufactured through the shop and when the production of the whole order is finished, they are moved to the stock point where the products are temporarily stored until they are requested by a customer. The customers require that their orders can be satisfied with a specified fill rate. Customer demands that cannot be met from stock are backordered. The production system consists of multiple functionally oriented workcenters through which a considerable number of different products can be produced. Each of these products can have a specific serial sequence of production steps, which results in a job shop routing structure. The production orders for different products compete for capacity at the different workcenters, where they are processed in order of arrival (FCFS priority).

To align the problem formulation with realistic situations, we assume that the customer demand process and the manufacturing process are subject to variability. We consider situations in which the average demand for end products is relatively high and stationary. Moreover, the production system is characterized by considerable set-up times and costs. In such a situation, companies are likely to apply batch-wise production. Typically, this type of production-inventory system can be found in metal or woodworking companies, e.g. the suppliers of the automotive or aircraft building industry. In particular, the advent of Vendor Managed Inventory (VMI) has forced manufacturing companies to integrate production and inventory decisions.

We assume that a single decision maker controls both the inventory points and the production system. We will present a decision procedure that allows us to determine production and inventory control decisions that minimize the total relevant costs in these make-to-stock PI systems. To this end, we will apply and integrate aspects from production control and inventory theory. Then, the objective of the optimisation procedure is to minimize the sum of set-up and ordering costs, work-in-process holding costs and final inventory holding costs. Moreover, we impose that the stochastic customer demand has to be satisfied with a prespecified fill rate. The decision variables that can be influenced by the decision maker are the review periods and the order-up-to levels of the products.

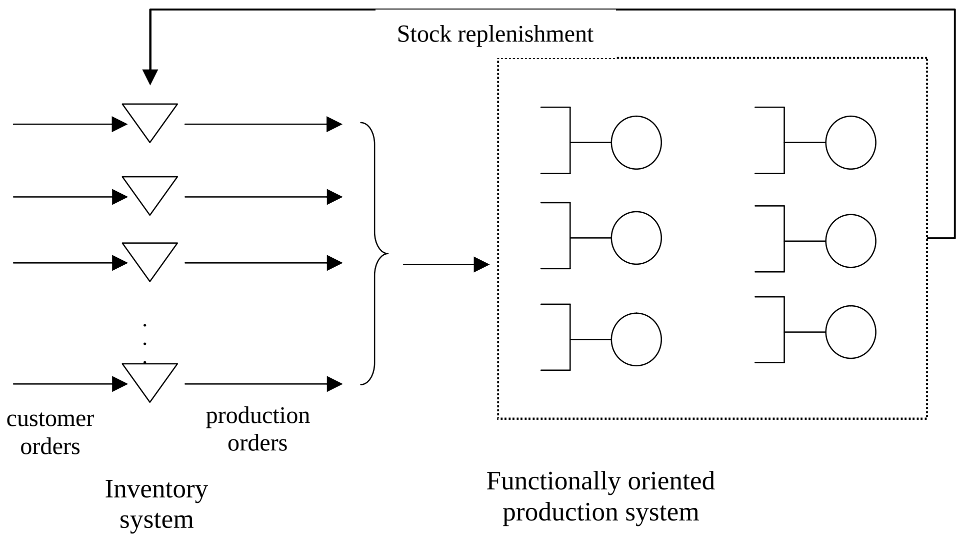

Figure 1: Integrated production-inventory system with job shop routings

To the best of our knowledge, no method exists that minimizes the total relevant costs in stochastic make-to-stock, functionally oriented shops. For related problems, however, solution methods have been developed. First, for the deterministic version of this problem, Ouenniche et al. (1998, 2001) present solution methods that are based on cyclical production plans. The cost-optimal cyclical plans are generated using mathematical programming techniques. Furthermore, for the case of a single workcenter PI system, the survey on the stochastic lot scheduling problem by Sox et al. (1999) gives an extensive overview of current research. Next, for the stochastic make-to-order job shop, Vandaele (1996) and Lambrecht et al. (1998) describe a lot sizing procedure that minimizes the expected lead times and thus, the expected work-in-process costs. To this end, they model the production environment as a general open queueing network. Finally, we mention the work of Bowman and Muckstadt (1995), who report on a production control algorithm for a stochastic multi-workcenter, multi-product production system that is cyclically scheduled. They use a Markov chain model to make a trade-off between inventory holding and overtime cost.

To minimize the total relevant costs in the stochastic make-to-stock, functionally oriented production system, we present a two-step decision making procedure that builds on the work of Vandaele (1996) and Lambrecht et al. (1998). The first step uses concepts from renewal and inventory theory, a queueing network analyser and local search techniques to set the review periods. This step uses several approximations, which may cause service levels to be lower than required. Therefore, we add a second step in which we fine-tune the order-up-to levels using a computational approach developed by Gudum and de Kok (2002). This approach uses the outcome of a simulation experiment to set the required order-up-to levels. Since our decision-making procedure uses approximations, we conduct an extensive simulation study to assess the estimation performance of the proposed procedure.

The remainder of this paper is organized as follows: section 2 gives an outline of the proposed decision-making procedure; section 3 presents the first step of the procedure in which the review periods and initial order-up-to levels are determined; section 4 discusses a method to fine-tune the order-up-to levels; in section 5 the procedure is tested in an extensive simulation study; finally, section 6 summarizes the findings of this research.

2. Outline of decision-making procedure

We briefly introduce the decision-making procedure that allows us to determine the control policy minimizing overall relevant costs in an integrated PI system. The cost minimization is constrained by a minimum required customer service level. The procedure builds on the observation that in an integrated PI system the replenishment orders that are generated by the stock points are identical to the production orders. Moreover, we assume that there are no transfer batches in the production system. Therefore, the only decision variables (on the tactical level) that can be influenced by the decision maker are the review intervals (R_i) and the order-up-to levels (S_i), for the different inventory points i, for i = 1, ... , P with P the total number of products. Finally, note that the review intervals determine the size and the arrival moment of production orders at the production system. By consequence, the arrivals of production orders to the production system, as well as the processing times of the production orders at the different workcenters are only determined by the review intervals and not by the order-up-to-levels. This implies that the delivery performance of the production system, e.g. the throughput times of the replenishment orders, is completely determined by the review intervals. The order-up-to levels only have an impact on the realized customer service levels.

We propose a decision-making procedure that consists of two major steps. In the first step, we model the PI system as a general open queueing network. We use parametric decomposition approximations (Whitt, 1983), elements from renewal theory (Cox, 1962) and standard inventory theory (Silver et al., 1998) to evaluate the total relevant cost of a given solution. Using a local search algorithm, we find the solution that minimizes the total relevant costs. This solution determines both the review intervals and the order-up-to levels.

However, we remark that the first step of our approach relies on approximations, which may have two serious consequences. Firstly, the solution found is not necessarily the optimal solution. Given the fact that no better queueing network analysing techniques are available at this moment, we think that it is not possible to develop a decision method that guarantees to find an optimal solution. Secondly, the initial solution can be an infeasible solution. The computation of safety stocks is based on approximations for several measures, which may result in service levels that are lower than required.

The second step of the decision-making procedure attempts to overcome the problem of infeasible solutions. In this step, we start by evaluating the performance of the initial solution (found in step 1) through a simulation experiment. The outcome of the simulation experiment is used to fine-tune the order-up-to levels so that the desired fill rates are achieved. In this fine-tuning step, we use a procedure developed by Gudum and de Kok (2002) that allows computing the required order-up-to levels based on the outcome of a single simulation experiment. This fine-tuning step builds on the observation that, for a given set of review intervals, changing the order-up-to levels only influences the realized customer service levels. This observation allows us to use the initial review intervals, fixed in the first step of the procedure, and adjust the order-up-to levels so that the customer service requirements are satisfied.

In the next sections, we will present the two steps of the decision-making procedure in more detail.

3. Step 1: setting review intervals and initial order-up-to levels

3.1. Formal problem statement

As stated in the introduction, the objective function considered in this paper is the minimization of total costs, which consist of set-up and replenishment costs, final inventory holding costs and work-in-process holding costs. First, we introduce some notation. Then, we derive formulas for the different cost components. After this, a formal problem definition will be given.

Notation:
P : number of products;
M : number of workcenters in the production system;
R_k : review interval of product k;
S_k : order-up-to level of product k;
A_k^C : interarrival times of customers for product k;
c_k : set-up and fixed replenishment costs incurred when placing one production order of product k;
v_{jk} : echelon value of one item of product k at workcenter j;
v_k : end value of product k;
r : inventory holding cost parameter;
$E[T_{jk}]$: average throughput time of production orders for product k at workcenter j;
ss_k : safety stock for product k;
a_k and a_k^* : realized and required fill rate for product k;

In a periodic review policy a replenishment order for product k is placed every R_k time units. Consequently, the number of orders placed per time unit is given by R_k^{-1}, so that the total ordering costs per time unit for product k are given by: $OC_k = c_k R_k^{-1}$.
The total work-in-process cost for product k is given by the formula below. We use Little's law to compute that the average number of items of product k at workcenter j equals $\frac{E[T_{jk}]}{E\left[A_k^C\right]}$. Multiplying the average number of products at a workcenter j with the echelon value and the inventory holding cost gives the work-in-process cost at a certain workcenter. Summing these costs over all workcenters gives the total work in process cost for product k: $WIPC_k = \sum_{j=1}^{M} \frac{E[T_{jk}]}{E\left[A_k^C\right]} v_{jk} r$.
The final inventory cost for product k is given by the formula below (Silver et al., 1998). The term between brackets gives the average amount of final inventory at stock point k, which consists of half the average order quantity plus the safety stock. Multiplying the average amount of final inventory with the end value of product k and the inventory holding cost gives the total final inventory cost of product k: $FIC_k = \left(\frac{R_k}{2E\left[A_k^C\right]} + ss_k \right) v_k r$.
The total cost for product k is simply the sum of its components: $TC_k = OC_k + WIPC_k + FIC_k$. Clearly, the total cost TC for the whole PI system is given by the sum over all products of the total costs for each product, so that $TC = \sum_{k=1}^{P} \left(OC_k + WIPC_k + FIC_k \right)$
Consequently, the mathematical formulation of the optimisation problem described in the introduction may be stated as:

$$\min \sum_{k=1}^{P} \left(c_k R_k^{-1} + \sum_{j=1}^{M} \frac{E[T_{jk}]}{E\left[A_k^C\right]} v_{jk} r + \left(\frac{R_k}{2E\left[A_k^C\right]} + ss_k \right) v_k r \right)$$

$$s.t.: a_k \geq a_k^* \qquad for\ k = 1,...,P$$

In the formal problem description, the terms $E\left[T_{jk}\right]$ and ss_k are unknown variables. In the next section, we will derive expressions for these unknown variables. These expressions will depend only on input parameters and on the review periods. In other terms, given a set of review periods we can compute the total cost of that solution. This property will be used to find the set of review periods that minimizes the total relevant costs.

3.2. Approximations for unknown variables

In this section we present a procedure to compute the unknown variables $E\left[T_{jk}\right]$ and ss_k, given a set of review intervals $R = \left(R_1,...,R_P\right)$. This procedure consists of three steps. In the first step, we determine the characteristics of the production orders that are generated by the stock points. In the second step, the throughput times through the production system are approximated using a queueing network analyser. Finally, the third step computes the required order-up-to levels and the corresponding safety stocks.

We start by analysing the generation of replenishment orders by the stock points. Note that in the PI system under study, the generation of a replenishment order at a stock point is equivalent to the placement of a production order to the production system. By analysing the characteristics of the replenishment orders, we therefore implicitly analyse the characteristics of the production orders that arrive to the production system, In our procedure, we focus on two main characteristics of the production orders: the time between the arrivals of two successive orders, referred to as the interarrival time A_k^P, and the processing time of the arriving production order for product k on workcenter j, denoted as P_{jk}^P. We confine ourselves to the determination of the average $E[.]$ and scv $c^2[.]$ of the interarrival times and the processing times. In the case of a (R,S)-inventory policy, a production order of variable size is placed at each review moment. Therefore, the average and scv of the interarrival times of production orders are given by: $E\left[A_k^P\right] = R_k$ and $c^2\left[A_k^P\right] = 0$. In a (R,S)-policy, production orders for product k are of variable size, which we denote here by N_k. By applying limiting results from renewal theory (Cox, 1962) we obtain that the number of arrivals in a review period R_k is approximately normally distributed with mean $E\left[N_k\right] = R_k E^{-1}\left[A_k^C\right]$ and variance $s^2\left[N_k\right] \approx R_k c^2\left[A_k^C\right] E^{-1}\left[A_k^C\right]$.

From these expressions we can derive the mean, variance and scv of the production order service times.

First, we introduce some extra notation.

P_{jk} : processing time of one item of product k at workcenter j;

L_{jk} : set-up time of production orders of product k at workcenter j;

Obviously, the average production time is given by the average total processing time plus the average set-up time, i.e. $E\left[P_{jk}^P\right] = R_k E^{-1}\left[A_k^C\right] E\left[P_{jk}\right] + E\left[L_{jk}\right]$.

We assume that the processing times of single units are i.i.d. and independent of the set-up time. Then, the variance of the net processing times, excluding set-up time, equals the variance of the sum of a variable number of variable processing times, which can be computed with a formula given by e.g. Silver et al. (1998). Consequently, the variance of the processing times of class-k orders at station j is given by:

$$s^2\left[P_{jk}^P\right] = E\left[N_k\right]s^2\left[P_{jk}^I\right] + E^2\left[P_{jk}^I\right]s^2\left[N_k\right] + s^2\left[L_{jk}\right]$$
$$= R_k E^{-1}\left[A_k^C\right]s^2\left[P_{jk}^I\right] + E^2\left[P_{jk}^I\right]R_k c^2\left[A_k^C\right]E^{-1}\left[A_k^C\right] + s^2\left[L_{jk}\right]$$

The second step in the procedure uses the characterization of the production orders to compute performance measures of the functionally oriented shop. Based on the characterization of the production orders, derived in the previous paragraph, we can model the production system as a general open queueing network in which the arrival and production processes of the orders have known first and second moments. From the late seventies on, extensive research has been done on the estimation of performance measures in such queueing systems. In our procedure, we use the queueing network analyser developed by Whitt (1983) in which we apply an improved expression for the scv of the departure processes (see Whitt, 1995). We refer the reader to the literature for more details on queueing network analysers, see e.g. Suri et al. (1993). The queueing network analyser allows us to find approximations for the average and variance of the throughput times at the different workcenters in the shop, i.e. $E\left[T_{jk}\right]$ and $s^2\left[T_{jk}\right]$. Assuming independence between workcenters, the average and variance of the throughput times can be summed to find the total replenishment lead time of a replenishment order for product k, whose average and variance are denoted as $E\left[T_k\right]$ and $s^2\left[T_k\right]$.

In the third step of the procedure, we determine the required order-up-to levels S_k using standard inventory theory. We require a characterisation of the customer demand D_k^{R+L} during the replenishment lead time and the review interval to determine the appropriate order-up-to levels, see Silver et al. (1998) for more details. The mean and variance of the customer demand during the order throughput time and review interval are given by:

$$E\left[D_k^{R+L}\right] = \left(E\left[T_k\right] + R_k\right)E^{-1}\left[A_k^C\right] \text{ and}$$
$$s^2\left[D_k^{R+L}\right] = \left(E\left[T_k\right] + R_k\right)s^2\left[\left(A_k^C\right)^{-1}\right] + E^{-2}\left[A_k^C\right]s^2\left[T_k + R_k\right].$$

Assuming that the demand during the replenishment lead time plus the review interval is normally distributed, the order-up-to level S_k can be determined by:

$$S_k = E\left[D_k^{R+L}\right] + k_k s\left[D_k^{R+L}\right]$$

where k_k is the so-called safety factor for product k that depends on the desired fill rate a_k^*, see e.g. Silver et al. (1998). The safety stock ss_k for product k can be computed as: $ss_k = k_k s\left[D_k^{R+L}\right]$.

In this section, we presented expressions to approximate the throughput times and the safety stocks, given a set of review periods. These expressions allow us to compute the total relevant cost of a given solution (see section 3.1). Based on the computational procedure presented above, the next section tries to find the set of review periods that minimizes the total relevant costs.

3.3. Optimisation

In the previous sections, we presented a procedure to compute the cost of a given set of review intervals $R = \left(R_1, ..., R_P\right)$. In this section, we try to find the set of review periods $R^* = \left(R_1^*, ..., R_P^*\right)$ that minimizes the total relevant costs. Unfortunately, we cannot prove the convexity of the objective function. Therefore, we resort to an optimisation technique for non-convex functions: simulated annealing, see e.g. Eglese (1990) for more details. We experimented with this method and compared its performance to a simple local search algorithm (searching a direction until no further improvement is possible). Surprisingly, there were no instances in which the simulated annealing algorithm performed better than the simple local search algorithm. Based on these experiments, we propose to use a simple local search algorithm in order to minimize the total cost. Again, we repeat the remark that, since our approach uses several approximations to compute the total cost of a solution, the optimum that is found using our procedure need not be the true optimum.

4. Step 2: computing order-up-to levels

In the first step of the decision procedure, we used approximations to determine the review intervals and initial settings for the order-up-to levels. Because of the use of approximations, it is possible that the initial solution that is computed in step 1 is an infeasible solution. The realized fill rates, e.g. observed in a simulation experiment, may be lower than the required fill rates. Also, it may happen that the realized fill rates are much higher than required. Obviously, this leads to a solution that is unnecessarily expensive.

For this reasons, we add a second step to the decision procedure. This second step starts by simulating the initial solution. This simulation experiment determines the realized fill rates. Based on the outcome of the simulation experiment, a procedure developed by Gudum and de Kok (2002) fine-tunes the order-up-to levels by setting them to the level that ensures the required fill rate, but also minimizes the final inventory holding cost. In other terms, the order-up-to levels are set to the lowest levels that satisfy the service level constraints. Note that we do not change the review periods that were computed in the first step of the decision procedure.

In this step, we use the property that the arrivals of production orders to the production system, as well as the processing times of the production orders at the different workcenters are only determined by the review intervals and not by the order-up-to-levels. Therefore, changing the order-up-to levels only influences the fillrates, without affecting the rest of the production-inventory system.

5. Testing the procedure in a simulation study

The presented procedure uses several approximations to determine the decision variables (review intervals and order-up-to levels). In this section we test the estimation performance of our approach through an extensive simulation study. We assess the quality of our procedure by comparing the cost estimates of the first step of our procedure to the outcome of a simulation experiment. Unfortunately, it is not possible to compare the solution of our procedure to the true optimal solution, since the true optimum cannot be determined.

In our simulation study, we consider an integrated PI system with 10 products and 5 workcenters. We assume that the customer demands arrive according to a Poisson process. Furthermore, the set-up times and processing times are exponentially distributed. This assumption allows incorporating all kinds of variability that are present in real production systems: operator influences, workcenter defects, etc. In the simulation study, we vary four factors over several levels:

- (i) net utilization of the workcenters (0.65, 0.75, 0.85);
- (ii) set-up times (randomly generated in the intervals [30, 60] and [90, 180]);
- (iii) set-up costs (randomly generated in the intervals [0, 0], [20, 40] and [60, 120]);
- (iv) fill rates (0.90, 0.98).

We generated three sets of random instances. Each of these instances is characterized by randomly selected routings, echelon and end values of products, demand rates, production times and inventory holding costs. The demand rates and production times are generated such that the net utilization of every workcenter equals the required level, while the production times are randomly generated from the interval [1, 5]. The routing structures are chosen so that the average number of operations per product equals 3 and the number of operations per product lies in the interval [2, 4]. Furthermore, the number of operations per workcenter lies in the interval [4, 8].

In this way, we generated three random instances. As mentioned, every instance has four factors that are varied over multiple levels. The number of experiments for one instance equals 3 x 2 x 3 x 2 = 36. Therefore, the total simulation study presented in this section consists of 3 x 36 = 108 simulation experiments.

Tables 5 summarizes the major findings of the simulation study. The table presents the relative difference between the performance prediction for the total relevant cost TC^{am} of our approximation model (described in section 3) and the total cost TC^{sim} observed in a simulation experiment. This relative error e can be computed as:

$$e = 100 \cdot \frac{(TC^{am} - TC^{sim})}{TC^{sim}}$$

The results are compared for the optimal solution found by our decision procedure. From table 5 it appears that the proposed approximation method performs rather well. The absolute error made by our approximations is 5 to 6 percent on average. The maximum error is up to 15 percent, which is still acceptable.

factors			**fill rate 90 %**			**fill rate 98 %**		
net utilization	**setup time**	**setup cost**	**route 1**	**route 2**	**route 3**	**route 1**	**route 2**	**route 3**
0.65	[30, 60]	0	5%	2%	7%	6%	3%	6%
0.65	[30, 60]	[20, 40]	5%	5%	6%	7%	6%	8%
0.65	[30, 60]	[60, 120]	5%	5%	6%	6%	7%	8%
0.65	[90, 180]	0	4%	4%	9%	5%	5%	11%
0.65	[90, 180]	[20, 40]	7%	6%	7%	8%	7%	9%
0.65	[90, 180]	[60, 120]	6%	5%	7%	8%	7%	8%
0.75	[30, 60]	0	-3%	-5%	0%	-2%	-3%	1%
0.75	[30, 60]	[20, 40]	6%	6%	8%	8%	7%	9%
0.75	[30, 60]	[60, 120]	7%	6%	8%	8%	7%	9%
0.75	[90, 180]	0	-3%	-5%	0%	-1%	-4%	2%
0.75	[90, 180]	[20, 40]	6%	4%	7%	6%	6%	8%
0.75	[90, 180]	[60, 120]	6%	5%	8%	7%	7%	10%
0.85	[30, 60]	0	-10%	-14%	-6%	-5%	-11%	-5%
0.85	[30, 60]	[20, 40]	2%	0%	4%	2%	0%	5%
0.85	[30, 60]	[60, 120]	3%	1%	5%	4%	2%	7%
0.85	[90, 180]	0	-8%	-15%	-6%	-7%	-12%	-3%
0.85	[90, 180]	[20, 40]	-5%	-7%	-2%	-5%	-7%	-1%
0.85	[90, 180]	[60, 120]	-1%	-2%	2%	0%	-3%	3%
		avg |e|	5%	5%	5%	5%	6%	6%
		min |e|	1%	0%	0%	0%	0%	1%
		max |e|	10%	15%	9%	8%	12%	11%

Table 5: Relative difference e between approximation method and simulation results

6. Conclusions

We propose a two-step procedure to make production and inventory control decisions in an integrated production-inventory system, characterized by job shop routings and stochastic demand, set-up and processing times. Our procedure tries to minimize total relevant costs (ordering and setup costs, work-in-process holding costs and final inventory holding costs) while stochastic customer demand is satisfied with a prespecified fill rate. The first step uses renewal theory, standard inventory theory, a queueing network analyser and local search techniques to determine the control parameters. Several approximations are used in this step. Since this may result in customer service levels that are too low or too high, the order-up-to levels are fine-tuned in the second step. This step ensures that all customer service level requirements are satisfied. We evaluated the estimation performance of the first step of our approach through an extensive simulation study consisting of 108 instances of the production-inventory system. The results of the simulation study indicate that our estimation method performs rather well.

Bibliography

Bowman, R.A. and J.A. Muckstadt (1995), Production control of cyclic schedules with demand and process variability, Production and Operations Management, Vol. 4, No. 2, p. 145-162.

Cox, D.R. (1962), Renewal theory, Methuen, London.

Eglese, R.W. (1990), Simulated annealing: a tool for operational research, European Journal of Operational Research, Vol. 46, p. 271-281.

Gudum, C. K. and T.G. de Kok (2002), A safety stock adjustment procedure to enable target service levels in simulation of generic inventory systems, working paper, BETA research school, The Netherlands.

Lambrecht, M.R., Ivens, P.L. and N.J. Vandaele (1998), ACLIPS: A capacity and lead time integrated procedure for scheduling, Management Science, Vol. 44, No. 11, p. 1548-1561.

Law, A.M. and W.D. Kelton (2000), Simulation modeling and analysis, Third edition, McGraw-Hill, Boston.

Ouenniche, J. and F. Boctor (1998), Sequencing, lot sizing and scheduling of several products in job shops: the common cycle approach, International Journal of Production Research, Vol. 36, No. 4, p. 1125-1140.

Ouenniche J. and J. W. M. Bertrand (2001), The finite horizon economic lot sizing problem in job shops: the multiple cycle approach, International Journal of Production Economics, Volume 74, p. 49-61.

Silver, E.A., D.F. Pyke and R. Peterson (1998), Inventory management and production planning and scheduling, Wiley, New York.

Sox C.R., P.L. Jackson, A. Bowman and J.A. Muckstadt (1999), A review of the stochastic lot scheduling problem, International Journal of Production Economics, Vol. 62, p. 181-200.

Suri, R., J. L. Sanders and M. Kamath (1993), Performance evaluation of production networks, Chapter 5 in Handbooks in Operations Research and Management Science, Vol. 4, S. C. Graves et al. (eds.), Elsevier, Amsterdam, p. 199-286.

Vandaele, N.J. (1996), The impact of lot sizing on queueing delays: multi product, multi machine models, Ph.D. thesis, Department of Applied Economics, Katholieke Universiteit Leuven, Belgium.

Whitt, W. (1983), The queueing network analyzer, Bell System Technical Journal, Vol. 62, No. 9, p. 2779-2815.

Whitt, W. (1995), Variability functions for parametric-decomposition approximations of queueing networks, Management Science, Vol. 41, No. 10, p. 1704-1715.