Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

BUFFERING AGAINST UNCERTAINTY IN HIGH-TECH SUPPLY CHAINS

Ton de Kok

School of Industrial Engineering
Eindhoven University of Technology
Eindhoven, AZ 5612, NETHERLANDS

ABSTRACT

In this paper we discuss the results of extensive research on supply chain modelling and analysis in high-tech supply chains. We distinguish between the forecast-driven and the customer-order-driven part of the supply chain. For each part we present a generic model that captures real-life complexity of high-tech supply chains. With each model comes a class of operational control policies that coordinate work order release across the supply chain. We discuss the empirical research that shows the validity of the models proposed and the associated class of control policies. We discuss the mathematical tractability of near-optimal control policies based on the fact that optimal control policies within the class of policies proposed satisfy generalized Newsvendor equations. We discuss optimal positioning of quantity and time buffers and issues for further research.

1 INTRODUCTION

High-Tech Supply Chains deliver user-specific capital goods assembled in a multitude of sequential and parallel steps executed by a multitude of different companies. Typically, these capital goods are delivered under volatile market conditions, while the cumulative supply chain lead time over which purchasing and manufacturing decisions are based on forecasts of the market demand ranges from 6 months to two years. In most situations the order-driven final assembly activities of the Original Equipment Manufacturer (OEM) adds only 10%-20% to the value cumulated from the forecast-driven purchasing and manufacturing activities. This implies that most of the capital investment in High-Tech supply chains is created on forecasts of the volatile future. As the customer/users of capital goods demand high on-time delivery this creates a high-risk environment to operate in.

There are a few levers companies in high-tech supply chains can use. Firstly, the volatile risky market situation implies that they should ensure high profit margins that enable to build a financial buffer against periods of low demand. Secondly, resource and pipeline alignment within the supply chain and with the customers/users enables to respond as quickly as possible to changing market conditions for the customer/users of the equipment produced by the OEM. The actual responsiveness of the supply chain to unexpected rates of demand increase depends on the buffers in resources and materials created in anticipation of such events. As anticipation buffers imply capital investment without an immediate return, the financial buffer needed in periods of low demand is partly consumed by anticipation buffers. In recent years we see an emergence of supply chain finance instruments being applied in high-tech supply chains, ranging from time-dependent trade-credit terms to customers/users investing in the OEM, and the OEM investing in key suppliers. The extend to which these instruments are used depends on the OEM market situation. In some markets the OEM is a *de facto* monopolist, in which case the supply chain finance instruments are used extensively. In other markets, such as health care equipment, competition is much stronger, whereby the OEM's position does not allow for favorable financing conditions to its suppliers, nor is it favoured financially by its customers.

In this paper we assume the market situation given. Eventually this translates into some market demand over time, which is partly known and partly unknown. The unknown part is considered to be stationary, i.e. exhibits a time-independent stochastic behaviour. We note here that this assumption of stationarity is a *statistically necessary* condition. If we would assume that the stochasticity is time-dependent, then each sample would have a unique underlying stochastic process, which cannot be inferred from this single data point. We also note here that the stationary stochastic behaviour can be translated into non-stationary stochastic behaviour, e.g. by assuming that the market demand exhibits multiplicative seasonal behaviour. Furthermore, the unknown demand may be uncertain in quantity and timing. Clearly, assuming that demand is forecast in future time buckets, one may abstract from the timing uncertainty. However, explicitly distinguishing between timing uncertainty and quantity uncertainty may improve the forecast quality. Still, in this paper we assume that we are confronted with either quantity uncertainty in discrete time or timing uncertainty in continuous time. Let us motivate this assumption, before using it in Section 2, where we propose a generic model for the high-tech supply chain.

When forecasting demand for the forecast-driven part of the supply chain, this typically concerns monthly or weekly demand forecasts. In many situations, weekly forecasts are derived by allocating the monthly forecast over 4 or 5 weeks according to some pre-specified percentages. The monthly and weekly buckets combined with the quantity volatility obscure possible demand timing uncertainty. Also note that demand for modules, subassy's and components delivered to the OEM by 1st tier suppliers is often derived from multiple types of equipment, which further obscures the demand timing uncertainty, which is fundamentally related to individual customer orders. Thus we assume that the forecast-driven part of the supply chain uses quantity buffers in the form of slack materials and slack resources.

In contrast, the order-driven part of the high-tech supply chain by definition is not subject to demand quantity uncertainty. Admittedly, we abstract here from possible changes in order specifications during the final assembly process. In principle these changes should be timed in such a way that they do not impact the final assembly process itself. This can be done by distinguishing between different phases in the final assembly process and communicating the last possible moment for changes in order specification elated to a particular phase. The timing uncertainty remaining is not related to the timing of order acceptance, but to the uncertainty in the final assembly process, particularly related to functional testing of modules, interfaces and the final product. Thus we assume that the order-driven part of the supply-chain uses time buffers in the timing of order release to the subsequent phases in the final assembly process.

The above distinction between the quantity-buffered forecast-driven part of the supply chain and the time-buffered order-driven part of the supply chain has been empirically tested in a series of MSc projects since 2002. The basis of the model for the forecast-driven part of the supply chain is the multi-item multi-echelon inventory model described in De Kok and Visschers (1999) and De Kok and Fransoo (2003). This model assumes constant planned lead times for each item, an arbitrary Bill of Material and i.i.d. demand in subsequent time units. The constant planned lead times imply that after an item work order release to the shopfloor, the item work order is completed after a fixed time, independent of the work order size. The multi-item multi-echelon control policy called Synchronized Base Stock policy ensures that work-order releases are material-feasible, so that only process uncertainty must be accounted for when setting the item-dependent planned lead times. We refer to De Kok and Visschers (1999) and De Kok and Fransoo (2003) for details on the policy structure. It is outside the scope of this paper to provide a detailed discussion of its mode of operation.

The basis of the model for the order-driven part of the supply chain is the multi-item multi-echelon production system model described in Yano (1987a), Yano (1987c), Yano (1987b), and Axsäter (2005). This model assumes a convergent (final) assembly process consisting of production stages, where each stage has a stochastic throughput time and throughput times of different stages are independent. The control policy assumed is based on constant planned lead times for each stage, which determines the planned release and completion times for each stage. A work-order for the final product consists of work-orders for each stage. A work-order for a stage can only start if all preceding work-orders have been completed and its start time has been passed. The model abstracts from capacity constraints and resulting congestion effects and solely takes into account the throughput time stochasticity caused by congestion effects. The model has been empirically tested at a manufacturer of lithographic systems for the semiconductor industry, showing that the model performance outcomes are close to the empirically measured performance outcomes, in particular the fraction of systems delivered on time to the customer (cf. Atan, De Kok, Van Boxel, and Janssen (2015)).

The contribution of this paper is as follows. Firstly, in Section 2 it presents a set of generic models and associated control policies that can be used as the basis for operational planning and control of high-tech supply chains. We discuss the empirical validity of these models as derived from the MSc projects mentioned above. The validation studies do not make any assumptions on the control policies actually used in practice. In fact, given the extensive human interventions, it would be impossible to describe the actual policies used. The ultimate consequence of the actual control policies used are the average inventory levels and average customer service levels measured over the validation period. Given that the Synchronized Base Stock policies assumed in the empirical studies, which are tuned to the average inventories measured, generate customer service levels in line with the actually measured customer service levels, we must conclude that the average inventory levels across the supply chain are the main determinants of the customer service levels. This is an extremely important finding, as the complexity of multi-item multi-echelon systems forces us to make assumptions concerning control policies that allow for computational tractability. Apparently we can make such assumptions without jeopardizing empirical validity. Secondly, assuming that the parameters of these control policies have been determined, either by algorithms or by human inputs, these control policies provide computationally highly efficient order release mechanisms that ensure all orders released are material feasible. This is clearly the case for the customer-order driven part. For the forecast-driven part of the supply chain, we note here that the computational complexity of the above-mentioned Synchronized Base Stock policies is similar to that of the MRP I logic (cf. Hopp and Spearman (1993). Thirdly, in Section 3 we consolidate earlier work (Diks and De Kok (1998), De Kok and Visschers (1999), De Kok and Fransoo (2003), Gong, De Kok, and Ding (1994), and Atan, De Kok, Van Boxel, and Janssen (2015)) to show that the optimal control policies satisfy so-called generalized Newsvendor equations. This enables efficient computation of near-optimal control policies. Finally, in Section 4 we discuss generic insights about where to buffer uncertainty, both in the forecast-driven and customer-order-driven part of the supply chain.

2 THE GENERIC SUPPLY CHAIN MODELS AND REALITY

The notion of supply *chain* is misleading. In real-life supply chains do not exist. We are faced with networks consisting of mutually dependent processing activities, coupled by controlled item stockpoints.

2.1 Forecast-Driven Supply Chain

At the most-downstream end of the forecast-driven part of the supply chain, we scope the supply chain by the products, the end-items, it delivers to a market place consisting of multiple customers, which demand, or more generally, the forecast error, is considered exogenous. Depending on the position of the customer order decoupling point (cf. Olhager (2003)), end-items are sellable products, subassemblies (modules) or components. This concerns Make-To-Stock (MTS), Assemble-To-Order (ATO), and Make-To-Order (MTO), respectively. In High-Tech Supply Chains ATO is the most common mode of operation.

At the most-upstream end of the forecast-driven supply chain we scope the supply chain by the suppliers, whose lead times of items supplied are considered exogenous. Within the scope of the supply chain we control item availability over time. Each item has a constant lead time that captures the inherent throughput times from order release to order completion that result from the confrontation between resources with stochastic processing times and stochastic item demand. The stochasticity in item demand is a consequence of the stochastic exogenous demand for end items that propagates upstream by item workorders released

Table 1: Model of the forecast-driven part of the supply chain.

i :	item index, where $i \in \mathscr{M} = \{1,...,M\}$.
k :	index of end-items, where $k \in \mathscr{N} = \{1,...,N\}$.
L_i :	planned lead time of item i.
$a_{i,j}$:	the number of items i used in each unit of item j.
h_i :	unit inventory holding cost for item i added to that of its child items.
H_i :	unit inventory holding cost for item i.
b_k :	unit backlogging cost for end-item k.
D_k :	demand for end-item k in an arbitrary period.
X_i :	net stock of item i at the end of period an arbitrary.
Q_i :	average order quantity of item i in an arbitrary period.

by the supply chain control policy. Despite the notion of dependent demand for upstream items, each time and again, for each item in the supply chain, the demand is stochastic. This fact makes multi-item multi-echelon systems so difficult to analyze mathematically, even under given control policies (cf. De Kok and Fransoo (2003)).

The model we propose for analysis of the forecast-driven part of the (high-tech) supply chain is presented in Table 1. We assume exogenous demand for end-items is i.i.d. in subsequent periods. We assume linear holding and penalty costs. In case penalty costs are unknown we can use the Newsvendor fractile to determine the penalty costs, i.e.

$$P\{X_k \geq 0\} = \tfrac{p_k}{p_k+H_k}. \tag{1}$$

Note that the Newsvendor fractile concerns the non-stockout probability at the end of an arbitrary period. This service measure is much easier to obtain from data in ERP systems than the often-proposed fill rate, which needs an accurate estimate of demand. In most real-life situations the notion of demand for end-items is far from obvious, e.g. in case of lost-sales. The matrix $(a_{i,j})$ is the Bill Of Material matrix that results from superposition of all end-item Bills of Material.

There is no hope that tractable optimal policies exist for multi-item multi-echelon systems (cf. Zipkin (2000) and Diks and De Kok (1998)) under stationary stochastic demand. In De Kok and Visschers (1999) and De Kok and Fransoo (2003) experimental evidence is provided that the Synchronized Base Stock policies perform well compared with other control policies for multi-item multi-echelon systems. Synchronized Base Stock policies are based on a generalization of the synchronization concept for pure assembly systems in Rosling (1989), which translates the multi-item multi-echelon network structure determined by the Bill Of Material and the planned lead times into a set of divergent so-called decision node structures. For such divergent structures Diks and De Kok (1999) proposes near-optimal policies composed of base stock policies for each decision node and linear allocation rules to ensure material feasibility of order releases in these divergent structures. The resulting order releases can be translated back to material feasible order releases in the original network.

The algorithms have been implemented in a software tool called ChainScope. As the tool has been applied in both high-tech supply chains (cf. Bisschop (2007), Goli (2005), De Jong (2010), Van Wanrooij (2012)), and in high-volume supply chains (Camp (2002), Den Hartog (2005), Hernandez Wesche (2012), Radstok (2013), Roose (2007)), we introduced the average order quantity Q_i as an input parameter. We implemented the average order quantity by setting the review period such that the average demand per review period was close to the average order quantity. This modelling approach was the only compromise to mathematical rigour. Both simulation experiments and our empirical studies showed that this modelling approach yields good results.

One may wonder whether the perspective of optimal policies is the only relevant one. We should always be aware of the fact that the concept of optimality is only rigorously defined within the context of a mathematical model with a unambiguously defined objective function. Any mathematical model is an abstraction of reality and for most models known in literature little evidence is provided of its empirical validity. Another perspective is that we need tractable models that explain the performance of real-life supply

Table 2: Model of the forecast-driven part of the supply chain.

s :	item index, where $s \in \mathscr{M} = \{0, \cdots, M\}$.
0 :	index of the customer specific end-item.
τ_s :	throughput time of stage s completing item s.
L_s :	planned lead time of stage s completing item s.
$a_{i,s}$:	1 if item u is the successor of s, 0 otherwise.
h_s :	inventory holding cost per time unit for item s added to the cumulative holding cost per unit time of its preceding items.
b_0 :	penalty cost per unit time late delivery of the customer-specific end-item 0.
X_s :	net stock of item s at the end of period an arbitrary.

chains. The MSc projects mentioned above and some more (Boulaksil (2005), Janssen (2004), Van Pelt (2015), Uquillas Andrade (2010)) showed that the Synchronized Base Stock policies and the algorithms derived from there, explain the relationship between on the one hand the structural input variables defined by the model, i.e. BOM, lead times, the historical data concerning exogenous demand, item inventories, and lot sizes, and on the other hand the historical data on end-item non-stockout probabilities and fill rates. We aggregated historical data over time into sample averages of item inventories, lot sizes and end-item customer service levels. From demand data we computed sample average and sample standard deviation of end-item demand.

Our empirical studies revealed another important phenomenon that has not been observed before in literature: multi-item multi-echelon control policies constrain the possible outcomes of average item inventories. This is a result of the mutual relationships between items via the BOM and lead times. Suppose two items have the same parent and item 1 has a longer lead time than item 2. In that case any sensible policy ensures that at any point in time that $X_2 \leq X_1$. It follows that $E[X_2] \leq E[X_1]$. In practice one may have decided to set target safety stocks of items 1 and 2 at the same level for "obvious" reasons. In that case e.g. MRP logic will yield $E[X_2] > E[X_1]$. The difference is *dead stock* of item 2. We implemented the dead stock concept briefly explained here as slack variables within the Synchronized Base Stock policies.

In conclusion, our empirical studies showed that the Synchronized Base Stock policies, together with the dead stock concept, could explain the empirically measured performance in a wide range of different multi-item multi-echelon systems that represented the forecast-driven part of real-life supply chains.

2.2 Customer-Order-Driven Supply Chain

By definition, the order-driven part of the (high-tech) supply chain does not face demand uncertainty. The uncertainty is primarily in the subsequent assembly steps and the functional testing of the resulting subassemblies and eventually the final product. Also here throughput times are the consequence of the confrontation of stochastic resource requirements in time and quantity, and resource availability. Again we use planned lead times for the various assembly stages distinguished to be controlled, to plan the order releases. As mentioned above we use the planned lead times to set planned release and completion times for each stage. Each stage starts at its planned release time or as early as possible after that, if it has to wait for completion of preceding stages.

In this case we explicitly use the probability distributions of the assembly stages as input for our optimization problem. Furthermore we assume that the cost added at a stage is incurred from the start of this stage until the product is delivered to the customer. Note that in this situation we assume that both holding costs and penalty costs are per unit time for each system assembled. With each stage we associate an item s, which is the output of the stage. By nature the final assembly process is strictly convergent. We note here that we assume that each customer-specific end-product follows the same stages. This is a necessary assumption, as otherwise we cannot plan. One may want to distinguish between several different product types, yet we need sufficient instances from the past in order to infer with some statistical support the stage throughput time probability distributions.

The model we propose for analysis of the order-driven part of the (high-tech) supply chain is presented in Table 2.

The throughput times τ_s for a single end-item are mutually independent. We do not need independence assumptions for the throughput times of different machines. We note here that in queueing network systems throughput times at the same stage are mutually dependent. This is not a problem. However, throughput times of subsequent stages may be dependent as well. The planned lead time control policy may induce independence of subsequent stages in case of a high probability of early completion of the predecessor stage. Furthermore stochasticity in throughput times in assembly stages may be dominated by the impact of testing and rework. In that case the independence assumption concerning different stages may be valid, as well. Clearly further research concerning the impact of the independence assumption is needed.

The objective is to minimize the sum of holding and penalty costs averaged over all machines. In Atan, De Kok, Van Boxel, and Janssen (2015) we use a conjecture that states that the optimal planned lead times solution satisfies generalized Newsvendor equations, which have a strong similarity to the generalized Newsvendor equations found in Diks and De Kok (1998). While the equations in Diks and De Kok (1998) can be solved recursively, the equations in Atan, De Kok, Van Boxel, and Janssen (2015) must be solved by an iterative fixed point algorithm, where in each iteration a set of generalized Newsvendor equations can be solved recursively. Thanks to the (experimentally determined) fast convergence of the fixed point algorithm, the optimal planned lead times solution can be efficiently computed.

In Atan, De Kok, Van Boxel, and Janssen (2015) we show the empirical validity of the model in the case of the final assembly of lithography systems for the semiconductor industry. We gathered empirical data on the throughput times, computed sample mean and standard deviation and fitted a gamma distribution. We gathered data on costs, planned lead times and on the sample fraction of systems delivered on time. With the given planned lead times we obtained probabilities of on-time delivery close to the sample fractions. Furthermore we found that in some cases the optimal solutions were close to the ones used in practice, in other cases substantial improvements could be made, i.e optimal overall planned throughput time for the final assembly process were shorter than used in practice. Further study revealed that the former situation occurred in situations were a considerable number of systems of the same type had been produced, while the latter situation occurred for newer systems with lower cumulative production numbers. We concluded that tacit knowledge, if sufficiently trained, can provide near-optimal solutions for assembly systems with stochastic throughput times. The model proposed supports much faster training.

3 GENERALIZED NEWSVENDOR EQUATIONS

As mentioned above, for both models proposed we can find the optimal solutions by solving so-called generalized Newsvendor equations. These generalized Newsvendor equations can be formulated by considering subsystems of the original system. Under the Synchronized Base Stock policies, the model for the forecast-driven part of the supply chain can be translated into a set of divergent systems that can be optimized independently. With each node n of a divergent system we associate a subsystem consisting of node n and all its downstream nodes. With each node n we associate an echelon holding cost h_n and cumulative holding cost H_n that can be easily deducted from the holding costs of the original system. With each node n we associate its downstream end-items $k \in E_n$. We define X_{kn} as the net stock of end-item k in the subsystem associated with node n. Then it is shown in Diks and De Kok (1998) that the optimal policy should satisfy the set of equations (2).

$$P\{X_{kn} \geq 0\} = \frac{p_k + H_n - h_n}{p_k + H_k}, \forall k \in E_n. \tag{2}$$

From the divergent structure of the subsystems it follows that equations (2) can be solved recursively. As the optimal policy is intractable, Diks and De Kok (1999) proposes to use so-called linear allocation rules. Unfortunately, these policies do not satisfy equations (2), but the solutions obtained perform well.

In Atan, De Kok, Van Boxel, and Janssen (2015) the generalized Newsvendor equations for the order-driven part of the supply chain were defined in a somewhat different way. They define the event A_s that a particular stage s causes late delivery of the final product to the customer:

An element w of the probability space is an element of A_s if and only if under w, i) stage s starts in time, ii) *if* the immediate successor of stage s starts in time then the final stage finishes in time and iii) the tardiness of stage s exceeds W_{-s}, and the final stage, stage 0, is late.

Here W_{-s} is the maximum of latenesses of all other predecessors of the successor of s than s itself. Informally, A_s represents the set of all events in which stage s is to be blamed for the lateness of the system. It is easy to see that the sets A_s, $\forall s \in \{1,2,\ldots,S\}$, are mutually exclusive and the union of these events consists of all w in the probability space for which the final stage is late. In this case the generalized Newsvendor equation conjecture is as follows:

$$P(A_s) = \tfrac{h_s}{p+H_0}, \forall s \in \{1,2,\ldots,S\}. \tag{3}$$

Now we can define an event B_n for the divergent system, similar to the event A_s for the assembly system. Note that we can associate with each end-item k a path from node s. With such a path we can associate a series of subsequent events, that consists of orders that start from an order by node s, subsequently an order at time $t+L_s$ from the immediate successor on the path to end-item k, and so on, until the order from end-item k, its subsequent arrival and the time of the next end-item k order arrival. Let us denote the time of this next order arrival as the critical moment. The event B_n is defined as follows:

An element w of the probability space is an element of B_n if and only if under w, i) the order of node n is satisfied by its predecessor, ii) *if* the order from the immediate successor of stage n on the path to end-item $k \in E_n$ is satisfied then there is no stock-out just before the critical moment and iii) there is a stockout at the critical moment.

Then equations (2) are equivalent to equations (4),

$$P(B_n) = \tfrac{h_n}{p_k+H_k}, \forall k \in E_n. \tag{4}$$

This implies that for both models considered, i.e. forecast-driven and customer-order driven, we find that the probability that an item (stage) causes disservice should be proportional to its added holding cost. As a corollary we find that the probability of disservice at an end-item equals its Newsvendor fractile. We believe that these results are very important in view of the computational efficiency that results from solving the generalized Newsvendor equations and the intuitive interpretation provided.

4 POSITIONING QUANTITY AND TIME BUFFERS IN THE SUPPLY CHAIN AND FURTHER RESEARCH

Given the above strong similarities between the optimization problems for the forecast-driven part and order-driven part of the supply chain, we expect similar findings concerning where to buffer against uncertainty under the (near-)optimal policies proposed. As mentioned in De Kok and Fransoo (2003) by far the most part of the supply chain buffers in time and quantity in the high-tech supply chain should be concentrated at the two main decoupling points. The first decoupling point decouples the forecast-driven supply chain from the order-driven supply chain by ensuring sufficient module availability to start the final assembly process in time. The second decoupling point decouples the process uncertainty in the subsequent (and parallel) assembly stages from the planned delivery moment by keeping a big time buffer at the last stage of the final assembly. This is in line with the conjecture in Goldratt (1997).

In the forecast-driven part of the supply chain other buffers are positioned immediately before assembly processes that produce items with a substantial added cost. In particular quantity buffers are created for predecessors of such items with long lead-times. In most of the supply chains studied in the MSc projects, the optimal policies do not hold any stock for many items, implying that the optimal policies create a flow from purchase items until the end-items. Though flow is strongly advocated by ”lean thinkers”, we found that most supply chain professionals consider our findings counterintuitive.

In the order-driven part of the supply chain we find that time buffers are held just before assembly stages, i.e. stages with multiple predecessors. As expected, the time buffers are highest for low-cost preceding stages. However, we found that for the lithography systems also time buffers were created for the most expensive module of the system. According to our intuition this is caused by the high throughput time uncertainty of this module. This reveals a specific aspect of the model for the order-driven supply chain that cannot be found in the model for the forecast-driven part: the impact of lead-time uncertainty.

This brings us to some issues for further research. Firstly, it is important to consider the two separate models proposed in an integrated setting. In high-tech supply chains, the planned lead times in the order-driven part of the supply chain have an impact on the forecast accuracy of the forecast-driven supply chain and the timing of the order releases in this part. The longer the planned lead times in the order-driven part to ensure on-time delivery to the customer, the earlier the orders for components and modules must be released and the lower the accuracy of the forecasts on which these order releases are based. In turn, the forecast-driven supply chain does not provide 100% fill rate, so that delays caused by disservice must be absorbed by the time buffers in the order-driven supply chain. As we argued that the added value of the modules constitutes already more than 80% of the final product value, the generalized Newsvendor equations seem to suggest that the modules should contribute considerably to the disservice to the final customer. That implies that the modules should cause substantial delays to the initial stages of the final assembly process.

ACKNOWLEDGMENTS

We acknowledge the efforts of numerous students from the School of Industrial Engineering at Eindhoven University of Technology who gathered empirical data at various companies, allowing us to test the empirical validity of Synchonized Base Stock policies and the software tool ChainScope based on these policies. We are grateful to these companies that provided the data and enriched the modelling of supply chains by their ways of working.

REFERENCES

Atan, Z., A. De Kok, R. Van Boxel, and F. Janssen. 2015. “Setting Planned Leadtimes in Customer-Order-Driven Assembly Systems”. Technical report, Eindhoven University of Technology.

Axsäter, S. 2005. “Planning Order Releases for an Assembly System with Random Operation Times”. *OR Spectrum* 27:459–470.

Bisschop, J. 2007. “Supply Chain Performance Evaluation: Application of the Synchronized Base Stock Policy in a High-tech Complex Equipment Supply Chain with Contract Manufacturers”. Master’s thesis, School of Industrial Engineering, Eidnhoven University of Technology.

Boulaksil, Y. 2005. “A Procedure to Set Safety Stock Levels in Multi-echelon Inventory Systems : A Rolling Horizon Simulation of the Supply Chains of Organon”. Master’s thesis, School of Industrial Engineering, Eindhoven University of Technology.

Camp, B. 2002. “Startrek Supply Chain Planning: Modeling, Optimization and Generalization”. Master’s thesis, School of Industrial Engineering, Eindhoven University of Technology.

De Jong, W. 2010. “New SCOP Method in ASML Supply Chain : Application of Enhanced Synchronized Base Stock in Planning of Supply Chain Planning Environment”. Master’s thesis, School of Industrial Engineering, Eindhoven University of Technology.

De Kok, A., and J. Fransoo. 2003. “Planning Supply Chain Operations: Definition and Comparison of Planning Concepts”. In *Supply Chain Management: Design, coordination and operation, Handbooks in Operations Research and Management Science*, edited by A. De Kok and S. Graves, Volume 11, Chapter 12, 597–675. Amsterdam: Elsevier.

De Kok, T. G., and J. W. Visschers. 1999. “Analysis of Assembly Systems with Service Level Constraints”. *International Journal of Production Economics* 59 (1): 313–326.

Den Hartog, B. 2005. "Tactical Supply Chain Performance Evaluation: Design of a Decision Support Tool". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Diks, E., and T. De Kok. 1998. "Optimal Control of a Divergent Multi-echelon Inventory System". *European Journal of Operational Research* 111 (1): 75–97.

Diks, E., and T. De Kok. 1999. "Computational Results for the Control of a Divergent N-echelon Inventory System". *International Journal of Production Economics* 59 (1-3): 327–336.

Goldratt, E. (Ed.) 1997. *Critical Chain*. USA: North River Press.

Goli, S. 2005. *Logistics Design Project at TBMS*. Ph. D. thesis, School of Industrial Engineering, Eidnhoven University of Technology.

Gong, L., T. De Kok, and J. Ding. 1994. "Optimal Leadtimes Planning in a Serial Production System". *Management Science* 40 (5): 629–632.

Hernandez Wesche, E. 2012. "Impacts of Implementing a Retailer Cross-dock on the Western Europe Procter & Gamble Supply Chain". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Hopp, W., and M. Spearman. 1993. "Setting Safety Leadtimes for Purchased Components in Assembly Systems". *IIE Transactions* 25 (2): 2–11.

Janssen, F. 2004. "Voorraadverlaging door SCM bij Diosynth". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Olhager, J. 2003. "Strategic Positioning of the Order Penetration Point". *International Journal of Production Economics* 85 (3): 319 – 329. Structuring and Planning Operations.

Radstok, K. 2013. "Fast & Slow Freight Distribution in the Fast Moving Consumer Goods Industry". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Roose, S. 2007. "Rethinking Inbound Operations Management at Procter & Gamble Mechelen: Multi-echelon Inventory Management Applied in Process Industry". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Rosling, K. 1989. "Optimal Inventory Policies for Assembly Systems Under Random Demand". *Operations Research* 37:565–579.

Uquillas Andrade, R. 2010. "An Integral Supply Chain Operations Planning System for a Global Pharmaceutical Company". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Van Pelt, T. 2015. "Multi-echelon Inventory Management at Sligro Food Group N.V.". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Van Wanrooij, M. 2012. "Strategic Supply Chain Planning in a Multi-echelon Environment: Identification of the CODP Location Constrained by Controllability and Service Requirements". Master's thesis, School of Industrial Engineering, Eindhoven University of Technology.

Yano, C. A. 1987a. "Setting Planned Leadtimes in Serial Production Systems with Tardiness Costs". *Management Science* 33:95–106.

Yano, C. A. 1987b. "Stochastic Leadtimes in Two-level Assembly Systems". *IIE Transactions* 19 (4): 371–378.

Yano, C. A. 1987c. "Stochastic Leadtimes in Two-level Distribution-type Networks". *Naval Research Logistics* 34 (6): 831–843.

Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill.

AUTHOR BIOGRAPHY

TON DE KOK is Professor of Quantitative Analysis of Operations at the School of Industrial Engineering at Eindhoven University of Technology in The Netherlands.He holds B.S. degrees in Mathematics and Physics, M.S degree in Matehematics from Leiden University and a Ph.D. degree from Free University in Amsterdam. He is director of the European Supply Chain Forum. His current research interest include

Build-To-Order Supply Chains and generic models for supply chain optimization under stochastic demand and supply. His email address is a.g.d.kok@tue.nl.