Production-Inventory Control Models: 40 Years of Relevant Research

Liber Amicorum for Ton de Kok

Editors:

Zümbül Atan

Willem van Jaarsveld

June 2025

Printed by: Grefo Prepress

Inhoudsopgave

Ton de Kok: Creatief, Energiek, Onverdroten

Een bijdrage leveren aan een Liber Amicorum voor iemand die je bijzonder graag mag en hoog hebt zitten is geen makkelijke taak, althans voor mij. Een ervaring die ik al had met een bijdrage die ik gevraagd was te leveren bij het afscheid van mijn leermeester Gijs de Leve aan de universiteit van Amsterdam. Als ik zwaar getafeld heb, droom ik er 's nachts nog wel eens over hoe mijn geestig bedoelde passages doodsloegen als bier in te vettig glas. Makkelijker vind ik het om in de pen te klimmen en mijn gal te spuien over de domheid en waanzin die onderwijsgevend Nederland van basisschool tot universiteit al jaren teistert en steeds erger lijkt te worden. Overigens zoden aan de dijk zet dit niet, maar het lucht wel op om tegengas te geven aan beurzensnijders en subsidie-slurpers rond het onderwijs.

Ton leerde ik kennen toen hij, op aanraden van Arie Hordijk, begin jaren tachtig solliciteerde op een vrijgekomen promotieplaats op het gebied van de stochastische besliskunde in mijn groep aan de Vrije Universiteit. De keuze voor Ton was niet moeilijk. Ik was meteen onder de indruk van de enthousiaste en energieke wijze waarop hij zich presenteerde in het sollicitatiegesprek. Als ik me goed herinner, had hij niet veel aan kansrekening en besliskunde gedaan aan de universiteit van Leiden maar des te meer aan fundamentele wiskunde, met name funktionaalanalyse. Een goede basis en doorzettingsvermogen gelardeerd met creativiteit zijn de beste ingrediënten die je hebben kan als je aan een serieus promotietraject in de stochastische besliskunde wilt beginnen. Zo voegde Ton zich op 1 april 1981 bij mijn besliskunde groep die destijds bestond uit de wetenschappelijk medewerker Frank van der Duyn Schouten – tegenwoordig de Dick Advocaat in het bestuursgilde van de rectores magnifici aan de Nederlandse universiteiten – de promotiemedewerker Michiel van Hoorn en de twee studentassistenten Luuk Seelen en Matthieu van der Heijden. Een fantastische groep waar Ton heel goed in paste, ook wat humor betreft. We hebben destijds wat afgelachen in de koffiekamer. Wie herinnert zich uit die tijd niet de Japanse gastbezoeker Toshi Kimura die bij een groepsdiscussie over wachttijdproblemen in zijn enthousiasme over een opkomend onderzoeksidee met blote hand sloeg op het tafeltje in de kamer van de op sabbatical zijnde decaan Nol Merkies. Dwars doormidden, nog veel moeite gekost het te repareren. In de groep werd niet alleen gelachen, maar ook hoogstaand onderzoek gedaan in de stochastische besliskunde. Dit bleek ook wel toen in 1983 een ranglijstje van onderzoeksgroepen in de stochastische operations research verscheen in het blad van de Applied Probability Society: de onderzoeksgroep aan de VU stond op de zevende plaats achter topuniversiteiten in Amerika en met binnen Europa alleen de groep aan Oxford university voor zich. Een feit dat mij goed van pas kwam toen in 1983 de universiteiten en hogescholen te maken kregen met de operatie Taakverdeling en Concentratie Wetenschappelijk Onderwijs (TVC) die tot doel had kleine studierichtingen – zoals de studie econometrie aan de VU – op te heffen of te concentreren en tevens het aantal hoogleraren fors te laten dalen. Zet dit laatste eens af tegen het huidige voornemen van sommigen om elke wetenschappelijke medewerker professor te noemen. Een variant van "geen gelul, iedereen een bul", een onvergetelijke uitspraak van Fred Steutel in zijn column ‘Effe zeuren’ van 26 mei 2011 in het informatie -en opinieblad Cursor van TU/e, waarin hij zijn verbazing uitspreekt over de tsunami van proefschriften aan de Nederlandse universiteiten en wat dit de universiteiten aan geld opbrengt.

Laat ik niet verder afdwalen en in mijn herinneringen duiken over de promotiejaren van Ton. In het eerste promotiejaar van Ton zette ik hem op een open ‘conjecture’, ik weet echt niet meer wat die conjecture behelsde maar ik weet nog wel dat het een doodlopende weg bleek waar Ton bijna een jaar lang met nietaflatende inzet op gezwoegd had. Na Ton’s eerste promotiejaar jaar besloten we het roer volledig om te gooien. Ton ging werken aan approximaties voor vooraadbeheersing en produktiebesturing. Een vulkaan van creativiteit barstte toen uit bij Ton met publicaties in vooraanstaande tijdschriften zoals Management Science en Advances in Applied Probability. In 1985 culmineerde dit in zijn proefschrift ‘Production-inventory control models; approximations and algorithms’. Na afronding van zijn proefschrift ging Ton werken bij het CQM, het consultancy bedrijf voor kwantitatieve methoden dat destijds onder leiding stond van Mynt Zijlsta. Voor zover ik weet was Ton de eerste van de VU afkomstige onderzoeker die bij CQM ging werken. Niet lang daarna gevolgd door mijn vroegere studentassistent Matthieu van der Heijden die later bij Ton en mij zou promoveren in het jaar 1992, ook het jaar waarin Ton tot voltijds hoogleraar aan TU/e werd benoemd. In de tussentijd was Ton ook deeltijds verbonden geweest aan de Universiteit van Tilburg waar hij een indrukwekkende reeks van memoranda schreef over approximaties voor wiskundige modellen in vervangingstheorie en voorraadtheorie. Deze exercities waren vingeroefeningen voor zijn latere fundamentele werk op het terrein van complexe netwerken van voorraadpunten die nauw met elkaar verbonden zijn en waar gegarandeerde service-niveaus essentieel zijn. Werk dat van groot praktisch belang is in de moderne high tech industrie. Ik heb dit werk slechts op grote afstand gevolgd, maar het is evident dat Ton als creatieve, praktische wiskundige uitgegroeid is tot een wereldtopper op het gebied van supply chain management. De internationale erkenning heeft even geduurd, maar Ton wordt tegenwoordig alom erkend te behoren tot de kleine selecte groep van leidende onderzoekers, waaronder Steven Graves van MIT en Hau Lee van Stanford University, op het gebied van supply chain management. De klasse van het werk van Ton en zijn onderzoeksgroep aan de TU/e blijkt ook uit de Edurank 2024 van de Best Universities for Industrial Engineering in the World, waar de TU/e op plaats 7 stond na topuniversiteiten als MIT en Georgia Institute of Technology, maar binnen Europa op plaats 1.

Het vervult me met enige trots en voldoening dat ik een steentje heb kunnen bijdragen aan de wetenschappelijk carrière van Ton. In de laatste paar jaar is Ton naast zijn werk aan de TU/e ook nog algemeen directeur (CEO) geweest van het befaamde Centrum voor Wiskunde en Informatica in Amsterdam. Het kan niet op. Ton, ik weet niet wat je na het bereiken van de pensioengerechtigde leeftijd gaat doen. Graag geef ik je nog het volgende advies. Ga een boek schrijven over al het werk wat je hebt gedaan op het gebied van voorraadbeheersing in supply chains waarmee je een brug hebt geslagen tussen wetenschap en praktijk. Je hebt voldoende materiaal liggen. Polijst dit tot een boek. Een boek heeft veel meer impact dan artikelen in tijdschriften: artikelen komen en gaan, maar een boek blijft bestaan.

Henk Tijms, Santpoort-Zuid, 10-12-2024

Filmscript, versie 0.0

Scene 1. Een zomerschool in een slaperig Engels stadje. Tijd voor de groepsfoto. De fotograaf houdt een afgerukte tak voor de camera om het beeld wat aan te kleden. Coryfeeën van toen en van nu, de dames vooraan. Midden achter kijkt een eerstejaars promovendus strak de lens in. Hij is direct door zijn promotor op reis gestuurd en valt op door zijn enthousiasme en weetgierigheid. Op het voetbalveld kun je hem beter niet tegenkomen.

Scene 2. Jaren later. De promovendus van weleer heeft een loopbaan in de industrie achter de rug. Hij praat met een collega die hij nog kent van de zomerschool, in een barak waarin een technische universiteit zijn afdelingen operations management en operations research huisvest. Het gaat over schedulingproblemen. Voor de één zijn dat vragen uit de praktijk die moeten worden opgelost. Voor de ander is het wiskunde, dingen met mooie en soms ook lelijke eigenschappen. Ze leren wat van elkaar.

Scene 3. Een warme middag in een promotiezaal. De promovenda is 's morgens in het huwelijk getreden. De verse echtgenoot en vier ouders zitten op de eerste rij, vader en schoonvader diep in slaap. Het wordt de promovenda teveel, ze valt flauw. De promotor springt op en vangt haar op, net op tijd. De echtgenoot denkt dat het erbij hoort en blijft zitten. De promotie gaat door, het begin van een mooie loopbaan in de consultancy.

Scene 4. Een vergaderzaaltje in de nog steeds niet afgebroken barak. Onderwerp: de oprichting van een onderzoekschool. De naam staat vast, vier letters. Heeft iemand door dat het een deelrij is uit de naam van de architect van de technische benadering van de bedrijfskunde die deze universiteit hoog in het vaandel schrijft? De scope staat ter discussie, want de erkenning is in eerste instantie uitgebleven. De oplossing: Beperk de deelname, benadruk de kwantitatieve aanpak en het theoretisch fundament. Less is more. Ook deze scene heeft een happy end. De aanpak maakt school.

Scene 5. Een congreshotel in de USA. Duizenden deelnemers en tientallen parallelsessies. De camera zoomt in op de delegatie van een operations management groep die nu een van de grootste ter wereld is en in de internationale top meedraait. Het werk van een charismatische en visionaire manager, die heldere doelen op de lange termijn stelde.

Scene 6. Twee hoogleraren begeleiden gezamenlijk een promovendus. Het project baart enige zorgen maar uiteindelijk toch ook een boekje. Meer vragen dan uitroeptekens. Voor de eerste promotor reden er een vervolgproject aan vast te knopen. De tweede promotor, inmiddels vertrokken uit de barak en zelfs van de universiteit, ziet later, tot zijn verrassing, zijn naam op een artikel dat uit dit achtjarig traject voortkomt. Het is zijn enige publicatie die hij niet heeft gelezen.

Scene 7. Een jaarlijks ritueel. Drie heren in een kaal vergaderzaaltje in het centrum van het land. Er ligt een stapeltje proefschriften op tafel. Welk daarvan krijgt de prijs van de onderzoekschool? De inzichten verschillen, de argumenten worden gewogen, de keuze is unaniem. Daarna praten ze over vroeger.

Scene 8. Jachtige muziek, grote stad. Een onderzoeksinstituut zit in zwaar weer. De nieuwe directeur toont een onbevangenheid die ruimte biedt aan originaliteit en vernieuwing. Het binden en stimuleren van onderzoekers, het combineren van wetenschappelijke kwaliteit en maatschappelijke urgentie, dat doet hij zijn hele leven al. Maar nu gaat het over algoritmiek en data, security en quantum computing, machine learning en artificial intelligence. Al is productie- en voorraadbeheer nooit ver weg.

Scene 9. Een heidag over een nationale strategie voor de wiskunde. Eén zegt: “We hebben meer geld nodig, knappe koppen voor de BV Nederland. En een centraal contactpunt voor innovatie, een brug tussen wiskunde en maatschappij.” Een ander: “Er is geen tekort aan geld maar aan leraren.” Een derde, een bekend gezicht: “In geen land ter wereld zijn de banden tussen wetenschap en bedrijfsleven zo hecht als in Nederland.” De wiskunde heeft altijd meer tijd nodig. Er wordt een nieuwe afspraak gemaakt.

Scene 10. Stilte. Een tuin, huiselijke gezelligheid. Een heer in een luie stoel, met een blocnote op schoot. Hij werkt aan stochastische probleempjes. Het beeld vervaagt.

Jan Karel Lenstra
20 augustus 2024

Travels with Ton: the early years

Henk Zijm[1]

Introduction

The invitation to write a contribution for a Liber Amicorum dedicated to my friend Ton de Kok inspired me to look back in particular to the period in which we operated as close colleagues within Philips, on our journey through the industrial landscape we were part of, and the scientific models and analysis we developed on our way. Traveling such an industrial road, of which about five years together, has been a rewarding experience for both of us, putting an undeniable mark on our scientific careers, and in hindsight I still feel grateful for the (relatively short) time, the early years, we were able to travel together. And although the word travel is here used in a metaphorical sense, we also literally travelled together every now and then; I still vividly recall our discussions with a beer at a terrace in Budapest, in between two conference meetings, or our joint singing of "Tulips from Amsterdam", one evening after another conference day, in a bar in Istanbul.

I am not sure about the first time Ton and I met, but it must have been somewhere halfway his period as a PhD student in Stochastic Operations Research at the Free University of Amsterdam, which lasted from 1981 till 1985. I myself had graduated at Eindhoven University of Technology almost four years earlier, on a topic in stochastic dynamic programming. During that time, I had become acquainted with most scholars of the Stochastic Operations Research community (which was quite close in these days), including professor Henk Tijms at the

[1]Emeritus professor of Industrial Engineering and former Rector Magnificus, University of Twente

Free University of Amsterdam, who later acted as Ton's supervisor. Ton's PhD research concerned inventory control, which I found interesting as I had meanwhile joined the Centre for Quantitative Methods (CQM), a consultancy office and at that time part of the Philips organization. There, I was confronted with many real life industrial problems, including inventories in worldwide supply chains, which I often discussed with peers in the field, including Ton. As a result, professor Tijms invited me to join the official graduation committee at the occasion of Ton's dissertation defense. I remember that Ton, uncharacteristic for him, appeared to be quite nervous. "What have I done to myself this time!" he exclaimed just a few hours before the start of the ceremony. Needless to say, the defense ran smoothly.

Shortly after the graduation ceremony, Ton shared his hesitation on what should be the next step. I do not remember all arguments but I recall that he, although certainly eligible for a continued career in academia, wondered whether almost ten years spent at universities wasn't enough, at least for the time being, and what other options were possible. And maybe I have suggested him, as I was enthusiastic about my own working environment, to write a letter to the CQM management, because with his background in inventory theory he might make a fine contribution to industry. Apparently, he did, as some weeks later an invitation had been sent to pay a visit to Philips-CQM for an interview.

State of the art in manufacturing and logistics

At this point, it may help to briefly sketch the state of the art of manufacturing and logistics planning and control within the Philips organization in the mid-eighties. It was the time that the MRP philosophy (Materials Requirements Planning and Manufacturing Resources Planning) did put a footprint in many industries, including Philips, as advocated by scholars such as Joseph Orlicky (Orlicky 1975), Oliver Wight (Wight 1981), and certainly also Clay Whybark (Vollman et al. 1984), a scientist we later frequently met, who passed away only recently. Although advocated as a revolution in production planning (Orlicky 1975), it soon became clear that MRP in its basic sense was primarily an administrative system, driven by a so-called Master Production Schedule (MPS) of given final product quantities needed in a number of future periods, which subsequently was exploded in a time-phased manner to lower levels of the product structure (Bill of Materials), through the use of off-set lead times (often expressed in terms of weeks). But the system did not answer important questions such as how to deal with uncertainty in either demand or supply, and how to calculate safety stocks or safety lead times. It lacked an intelligent capacity or resources planning system, or a workload control concept, topics later addressed in the Netherlands by e.g. Bertrand et al. (1990). And although more advanced concepts were available at that time, cf. Hadley (1963), Hees and Monhemius (1972), Peterson and Silver (1979), they were often considered to be "too theoretical to be applicable". On the other hand, it should be said that many OR scientists were not prepared to address the complexities of real world industrial problems, but merely concentrated on mathematical solutions for relatively small, isolated problems. So, yes, a significant gap between theory and practice existed, even amplified by notions, e.g. by Orlicky (1975), stating that "such mathematical approaches were no longer needed in view of the computing power of modern computers that would solve all planning issues" (sic!).

Long lead times and corresponding high work-in-process inventories were another characteristic of many industries in the late seventies. This could largely be attributed to the strong departmental structures and separation of concerns that existed in plants and factories in these days. Production efficiency used to be the guiding control principle, leading to large batch production in every production phase, in particular in mass production. This effect was even amplified by the offshoring of production to low-wage countries in the Far East, again inducing long supply lead times in large batches to efficiently use transport means. Figure 1 shows a typical supply chain (although we did not use that term very often in the mid-eighties, we merely spoke of worldwide logistic chains). But what one typically observes is: stocks everywhere around the globe, which in addition require enormous investments in warehousing and storage capacities, while insurance and also obsolescence costs were often at stake. The pressure was to reduce stocks and associated lead times wherever possible.

That was the situation within not just Philips but in many industries in the early eighties. Naturally, the Board of Philips put quite some pressure on finding remedies to cope with these problems. In our view, a systematic, mathematically based approach might offer one of the tools to address the situation. That was the world we were living in, and which Ton entered in late 1985.

Intermezzo

Here I have to make a small intermezzo: that entrance was not entirely smoothly. Ton traveled to Eindhoven, subsequently made his way to the end of the Boschdijk where the headquarters of Philips and also the CQM office were located at that time, and observed that a large iron fence was surrounding the company territory. As he was a well-trained mathematician, he quickly determined the shortest path,

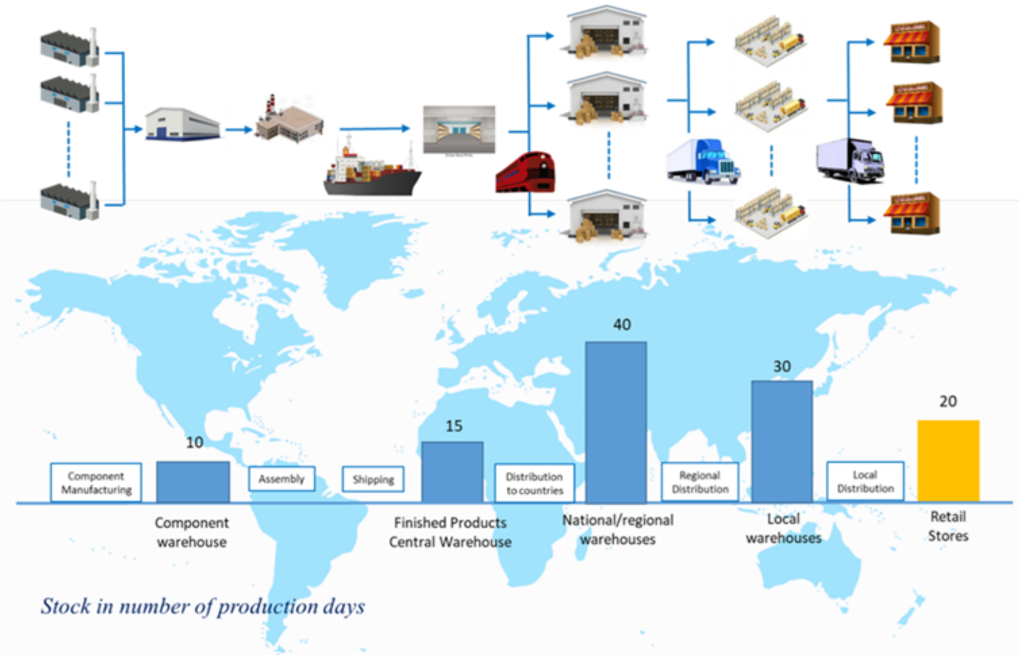

Figure 1: Global supply chain for consumer products

and perhaps he had already studied Goldratt's theory (Goldratt 1980), stating that any unnecessary obstacle might be removed, so he quickly determined that climbing the fence was the fastest way to arrive at his final destination. Shortly after that, the CQM management was notified by security that someone had made a somewhat peculiar entrance. Fortunately, the interview was not abandoned, otherwise a splendid career in both industry and science might have been stopped even before it started.

Collaborating in industry

The positive decision of the CQM board marked Ton's entrance into the world of industry, and we became colleagues during the next five years or so, until I left Philips. I can only say these five years formed a great time, in particular when collaborating in projects (which happened a few times), but also because we found time to discuss the essentials of problems we encountered, essentials that could not always be addressed in the course of a project but which lend themselves for a more thorough theoretical analysis, or for example for a master's student thesis topic, under our supervision. But we also learned that convincing practitioners at the shopfloor of our ideas was not always easy, to put it mildly. I mentioned already the in our view important shortcomings of MRP. Another example concerned our plea to install buffers in production lines, after which we were told that we were too much absorbed by our math analysis and that we seemed to have missed the revolution towards “Zero stock manufacturing” that was going on in industry. It should be said that at the same time the confrontation with new developments in industry was a great learning experience, on the paradigm shift towards quality and flexibility, see e.g. Schonberger (1982) and, within Philips, on the nature of the various product divisions Philips was composed of in these days, including the distinction between consumer-oriented (B2C) and business-oriented (B2B) divisions. In particular the work of Hoekstra and Romme (1986) on Customer Order Decoupling Point (CODP) principles was useful in structuring control policies depending on product-market characteristics.

Naturally, we learned most from the various projects within Philips almost always starting with a question from a production or logistics manager in whatever product division. One of these projects in which we collaborated, I will discuss in some more detail, because it illustrates various aspects of a large industry practice in these days. When looking back at the results of that particular project with today’s eyes, the insights we gained may seem trivial, but that’s an observation in hindsight. The study concerned the logistics of a new sort of telecommunication exchange, the SOPHO-S exchange for voice- and data-transmission, at Philips Telecommunication in Huizen, close to Amsterdam (cf. Zijm and De Kok 1989). Such exchanges were to be installed in large national telecom organizations in a variety of countries but also in private businesses worldwide. It concerned a typical professional market, meaning that the final assembly of the exchanges, that were composed of a number of cabinets each filled with cables and many (large and expensive) Printed Circuit Boards (PCB’s), took place only once the customer order had been firmly established, i.e. final assembly was production-to-order, or after the customer order decoupling point. Whereas the many components and subassemblies needed to build the final exchange were often produced based on forecasts of what might be needed, i.e. components and subassemblies fall under a production-to-stock regime, hence before the customer order decoupling point (cf. Figure 2).

In particular, the PCB assembly was based on forecasts, hence with all the uncertainties around these forecasts, which therefore require safety stocks to avoid a lack of required subassemblies or components once the customer order had been established. Unfortunately, each different destination country used different transmission functions which were reflected in specific designs of quite a number of PCB’s. In addition, the production lead time of these PCB’s might be up to

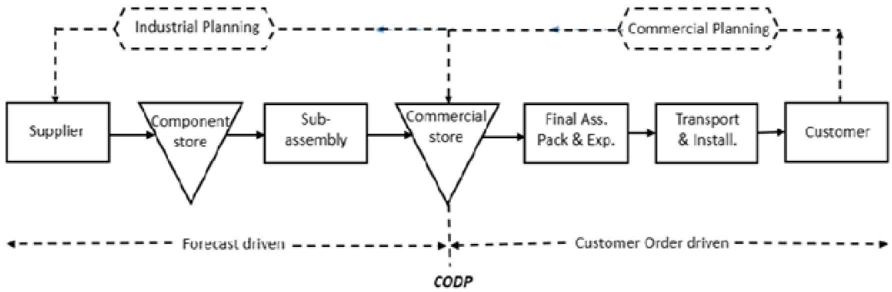

Figure 2: Logistics diagram of the Exchange production process

five months (for a variety of reasons that are beyond the scope of this paper). Using a fairly simple (R, S_t) inventory model, we were able to demonstrate that a significant reduction of the PCB assembly time might obviously dramatically reduce the inventories of assembled PCB's, see Figure 3. In our view, an even more important notion was a further reduction of safety stocks by a change in product design, leading to a severe reduction of product variety at the CODP. The idea was trivial. Produce only generic PCB's to be used for any destination country in the forecast-driven phases, whereas the country-specific transmission part, that represented only a very small part of the entire PCB's, might be constructed and added as a plug-in to the generic PCB, only after the customer order was established. Clearly, any undergraduate student in stochastic OR knows that product variety reduction with the same total forecasted number may reduce variability. Another way to phrase this in today's terminology (but no part of the literature at that time), is design for postponement, or modularity in product design, where customer-specific parts are added in a sort of plug-and-play assembly. However, trivial as it may sound, the engineers were not very responsive. They raised a number of technical arguments why such a modular design was not possible, or at least would induce additional assembly costs (which was true). But what essentially dominated the debate is that design engineers and logistic engineers in these days lived in two different worlds, i.e. the strong departmental structure that characterized not only Philips but the far majority of industries at that time prevented integral thinking, let alone integral logistics thinking. And, although trivial, even today we still observe many examples of a situation in which local optimization may severely hamper overall improvements, in industry but also in for example environmental issues or the energy transition. The adage "their costs are not our costs" continues to prevail.

Looking back, I still feel that the telecommunication project was a great learning experience, as was in fact the case for many projects. Another joint project that I

Action	Reduction of SA Inv. Pos.
Reduction of Subassembly lead time	
5 months → 3 months	26%
3 months → 1 month	36%
5 months → 1 month	52%
Advanced demand infomation	
Pre-inform. Order arrival	24%
Pre-inform. Order arrival and size	36%
Higher degree of commonality	
Country-independent PCB's	10%

Combination of all actions yields a maximum possible reduction of the inventory positions of 72%.

Figure 3: Percentual reduction of assembly stock values

briefly like to share, in particular to highlight Ton's feeling for exposure, concerned the assembly of coffee machines in a plant in Groningen, part of the Division of Small Domestic Appliances and Personal Care within the Philips organization. Imagine an accumulating conveyor system which is used to transport work pieces, fixed on pallets or product carriers, that have to visit a number of workstations such that at each workstation some essential components are mounted on what eventually becomes the coffee machine. Routings are job-dependent, hence not every job necessarily visits each workstation (cf. Figure 4). That gives rise to a rather familiar assembly line layout as you can see in the picture. We developed some control algorithms for that specific structure but what I remember in particular is Ton's presentation of the results. He called the control policy the "crenellation routing" (in Dutch: kantelenroutering), because it mimics the shape of battlement fortifications, and when asked why attributing such a fancy name to something so obvious, his response was: "Well, if you want to be successful, it helps to profile yourself a bit". I admit: it worked.

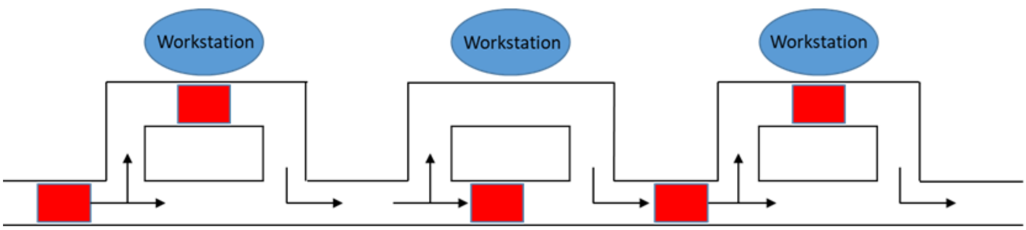

Figure 4: The coffee machine assembly line architecture, or crenellation routing (cf. Ton de Kok)

From industry to science

As noted earlier, we both liked to translate the problems that we encountered in factories into a mathematical model, and not only to discuss it at the blackboard in our offices but subsequently to give it to a student or trainee at CQM; in fact we supervised quite a number of master students who liked to study a realistic industrial problem as part of their final thesis work. One example of such a thesis, of Michiel Klaren of the University of Twente, and supervised by both Ton and me, was the buffer architecture in production lines. Imagine a set of workstations arranged in a line architecture, where each workpiece needs to visit each workstation, and all in the same order (Figure 5). If however the processing time at each workstation may vary (is stochastic) or if workstations may breakdown once in a while, the failure of one such workstation may induce the entire line to stop; downstream stations are no longer fed whereas upstream stations are blocked. To increase productivity, or to avoid a complete line breakdown due to the failure of a single workstation, it is wise to install some buffer places in front of each workstation, so that the non-failing workstations may continue working. But of course, buffers shouldn't be too large because they require space, and their occupation adds to the total amount of work in process, as well as the total lead time needed to finish an entire job. So, what's an appropriate size of the buffers. That problem had been studied already in the early seventies, by e.g. John Buzacott (1972) , and here in Eindhoven e.g. by Jacob Wijngaard (1979) and his PhD René de Koster (1988), now professor at Erasmus University. Ton was definitely inspired by the problem has he later wrote several theoretical studies on buffers in production lines and obtained a number of fundamental results (see e.g. De Kok 1990a).

Figure 5: Allocation of buffer space in flow line

A topic that has absorbed a lot of attention of both of us is that of multi-echelon systems, i.e. systems composed of what we now call a complete supply chain, starting with the delivery of raw materials and ending with final products stored at sales outlets, at least if they concern consumer products, and all aspects in between. Such supply chains were really at the heart of many product divisions of Philips, as were the corresponding problems of long lead times and high Work in Process inventories, as outlined earlier. As I meanwhile had become a parttime professor at the math department of the Eindhoven University of Technology, next to my job in CQM, I concentrated on more theoretical aspects, among others with

Geert-Jan van Houtum as a master student. While Ton, still fully employed by Philips, started to work with engineers in the various product divisions to convey the ideas of multi-echelon inventory theory in a supply chain framework. So, although the topic coined the interests of both of us, the directions we took were different. Geert-Jan, first as a student and later as my colleague in Twente, and I produced a series of papers, on pure cost models, hence models characterized by inventory holding cost and penalty costs in case of delivery failure, extending in particular work of Clark and Scarf (1960) and Eppen and Schrage (1981) on supply chains (see e.g. van Houtum and Zijm 1991). A right criticism, not only by Ton but many others as well, was that penalty costs may be somewhat arbitrary, and although it is easily shown that each combination of holding and penalty costs had a one-to-one translation to holding costs and service levels (Houtum and Zijm 2000), it of course remains an indirect approach. Ton felt more comfortable, and certainly in his more elaborate contacts with production and supply chain engineers in Philips, in taking a direct service perspective, and that has brought him considerable successes, both scientifically, cf. Verrijdt and De Kok (1996), Diks et al. (1996), Diks and De Kok (1998), Diks and De Kok (1999), De Kok and Visschers (1999), but also in his work with several product divisions, in particular in Philips Consumer Electronics and Philips Semiconductors.

Separate routes

As I said earlier, in 1990 our roads split, as I left both CQM and the math department of Eindhoven University, and not long after that, Ton was invited to fill the place that I had left, as a parttime professor in the math department again, which lasted until he became a full professor in the Department of Industrial Engineering and Management Science of Eindhoven University. As I started to work as a full professor in the Department of Production Engineering in Twente, I concentrated on smart manufacturing and maintenance service logistics, while Ton continued his work on service-oriented multi-echelon systems in which he became a recognized expert. But we continued to keep in touch, both became fellows of the BETA research school, and discussed problems of mutual interests. And of course, if a PhD was near completion, we invited each other a few times in the respective graduation committees. I spent slightly more than one year back in Eindhoven, in which year we also had one joint PhD on multi-echelon systems in Service Logistics, Jan Willem Rustenburg, a Naval Officer at that time, with Geert-Jan as daily supervisor (cf. Rustenburg et al. 2000). But, as my family lived in the Twente region, and daily commuting was not an option, I returned to Twente, where I quickly got involved in a number of university governance jobs which prevented me from doing much research at all. Ton made a different, and anyway more consistent choice, staying in science albeit with a strong eye on industrial relevance, and building a network of colleagues and groups worldwide, e.g. in the European Supply Chain Forum which he led for many years.

Interactions at Dinalog

But it turned out not to be the end of our interactions. After having served in various management jobs at my university, I decided in 2009 to finally return to my academic loves: education and research. Shortly after that, I received an invitation to also accept a parttime job as the Scientific Director of the Dutch Institute for Advanced Logistics (Dinalog), in Breda, somewhat in between Rotterdam and Eindhoven. Dinalog had been founded as the result of an initiative of the Dutch Ministry of Economic Affairs and the Ministry of Infrastructure and Water Management, to install a number of focus institutes in sectors that were considered to be crucial for the Dutch economy, and Logistics was rightly labeled as one of them (although it took some lobby work). The idea was that projects should be launched and executed based on a strong collaboration of academic and industrial partners, partly funded by the Dinalog (public-private partnerships). As both Ton and I had a keen interest in scientific research with a strong industrial or societal relevance, it became natural that we met again, although absorbed in clearly different roles. My role was to select proposed projects and, once approved, to make sure that part of the funding was allocated by DINALOG, whereas the companies and institutes involved might contribute the remaining part, either in kind or in cash. Ton played a leading role in one of the most important projects executed, to be discussed below.

A serious problem at the start of Dinalog was that some industrial partners narrowed the concept of logistics to multi-modal transport and warehousing, whereas many academics felt the topic of logistics should even be broadened to Supply Chain Planning and Control. And indeed, the latter became an important working field, next to subjects such as Service Logistics, Supply Chain Finance, multimodal corridors, transnational border management, and others. Naturally, Ton was at the heart of the Supply Chain Planning and Control field, active in various projects of which I mention here in particular the 4C4More project on Fast Moving Consumer Goods, working together with key researchers in related fields, e.g. Jos van Hillegersberg in Business Information Systems, my colleague at Twente University, and with representatives of companies such as Unilever, Kuehne+Nagel and others.

Figure 6: Transport, logistics and supply chains

To briefly discuss the 4C4More project and Ton's role in that, let me start with the definition of 4C. That term was initially coined by another old friend, Peter van Laarhoven, who authored a key report on the Logistics Innovation Program for the Ministry of Infrastructure and Water Management (Laarhoven 2008), which laid the foundation for what would become the institute Dinalog. 4C initially stood for Cross Chain Control Center (hence the four C's) but as some people disliked the word "control" or considered it too much reflecting a top-down approach, they phrase it as Cross Chain Collaboration Center (by means of a control tower!). The basic idea is that supply chains are seldomly controlled in its entirety by one entity, one firm say, but even more important, that various supply chains may exploit similar resources. Hence, it seems wise to not only integrate planning and management alongside a supply chain but also across various supply chains, to improve scale benefits, provide access to new channels, improve cost efficiency and diminish environmental burdens.

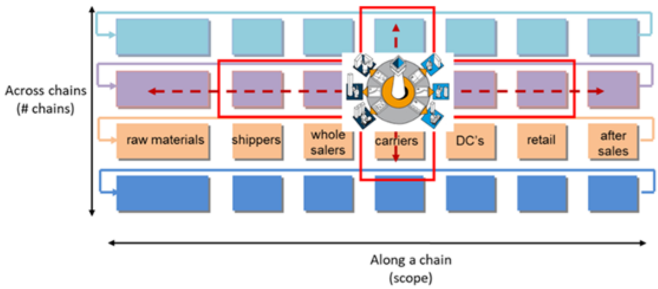

Figure 7: Coordination along and across supply chains by a control tower

The 4C4More project is a splendid example of what can potentially be achieved through such a coordination in two directions, in this case in particular in the Fast Moving Consumer Goods sector (De Kok et al. 2015). To set the record straight, I was not involved in the 4C4More project as such; I only had become instrumental in providing part of the necessary funding, and in monitoring the progress of sub-projects and operational objectives, given financial (funding) constraints.

One more activity, and herewith I complete this personal account, was the organization of a conference on Logistics and Supply Chain Management in Breda, the city where the Dinalog office was located. Again, we were able to closely work together in setting up the conference, which was partly hosted by the Military Defense Academy in Breda with which Dinalog had established firm contacts. Shortly after that, I left Dinalog to finally get back to one job at the time, my full chair at the University of Twente, and Ton concentrated on his research activities in Eindhoven again.

A final note: Dear Ton, the years of our collaboration (our joint travel) have not been forgotten, as witnessed above. I consider them as an indispensable part of the journey through my professional life. Thank you for your friendship: I wish both you and Irene a great time ahead.

Van Medewerker tot Baas

Mynt Zijlstra, oud-baas bij CQM

Inleiding

Naarmate je ouder wordt krijg je meer zicht op hoe je geheugen werkt, maar leer je het steeds minder begrijpen.

Spontane herinneringen aan toen Ton de Kok bij CQM werkte (1985 – 1990) te over, maar geen van alle betreffen de inhoudelijke bijdragen die hij geleverd heeft. En dat moet toch zijn belangrijkste tijdsbesteding geweest zijn, mag je aannemen. Gelukkig zijn er publicaties uit zijn CQM-periode en als je die herleest komt er uit de krochten van je geheugen wel weer wat bovendrijven.

Eerst een paar spontane herinneringen en daarna meer over het echte werk.

Spontane herinneringen

Aan het begin van Ton's eerste werkdag bij CQM (toen nog een afdeling van Philips) moest hij zich melden bij de Personeelsdienst in de Willemstraat. Vervolgens naar CQM, gevestigd op complex Vredenoord. Daar was vanwege bezuinigingen vanaf 9.00 uur de poort aan de Pieter Zeemanstraat dicht. Nou zijn in Tons leven barrières alleen maar uitdagingen. Hij klom over het hek aan de Boschdijk.

Al ras was Ton met zijn anekdotes en oneliners populair bij de collega's. In de kantine van VN zat half CQM (met de korte zijde van het dienblad naar voren!) bij Ton aan tafel. Er werd veel en heel luid gelachen. En dat was niet wat de serieuze ISA-bewoners van VN van de CQM-nerds verwachtten.

In één van zijn oneliners typeerde Ton zijn toenmalige collega Henk Zijm als “ome Henk, de bard van Diemen”. Jarenlang bleef het een gevleugelde uitdrukking binnen CQM.

Ton was een getalenteerde amateurvoetballer. Hij speelde in Best bij Wilhelmina. Of was het bij Best Vooruit? Hoe het ook zij, een aanwinst voor het CQM-elftal bij de ISA-voetbaltoernooien. De foto’s laten zien dat Ton vooral in de rust excelleerde.

In het voorjaar van 1990 lag ik met een aantal gekneusde ribben ziek te bed. Met name lachen was heel erg pijnlijk. Waarom CQM juist Ton de Kok had aangewezen om mij een ziekenbezoek te brengen? Toen hij de door mij bedachte touwconstructie om in en uit bed te komen zag, sprak hij: “Nooit geweten dat jullie zo’n SM-stel waren”.

Ton was altijd en overal bezig met zijn vak. In zijn kantoor en zelfs bij hem thuis struikelde je over de briefjes met ideeën en berekeningen. Bouwstenen voor onderzoeksresultaten. Dat brengt ons bij het echte werk.

Ton demonstreert aan Guus van Dongen
hoe je met een omhaal een doelpunt maakt.

Tactisch overleg
in de rust.

Het echte werk

Opvallend vaak zie je dat ervaringen opgedaan tijdens slechts enkele jaren werken in de praktijk een grote rol spelen in het oeuvre van academici.

In de 50-er en 60-er jaren gingen wiskundigen praktijksituaties (wachtrijen, voorraadregels e.d.) bestuderen. Ze stileerden de problemen en kozen hun aannames zodanig dat een mooi stukje exacte wiskunde ontwikkeld kon worden.

Helaas zijn praktijksituaties complexer en probleemeigenaren herkenden hun werkelijkheid niet in de bestudeerde modellen. Dat kwam het imago van de Toegepaste Wiskunde en meer in het bijzonder de Operations Research niet ten goede.

Ton de Kok zegt er in zijn inaugurele rede, uitgesproken op 22 januari 1993 aan de TU/e, over:

Het werken met mathematisch correcte resultaten biedt zekerheid, werken met benaderingen vergt durf, gebaseerd op een zeer gedegen mathematische ondergrond. De complexiteit van de modelanalyse is echter niet relevant voor de probleemeigenaar. Wel het model, waarvan moet worden nagegaan of het inderdaad een adequate beschrijving geeft van de werkelijkheid en de ervaren problematiek.

Het analyseren van zulke door de probleemeigenaar geverifieerde modellen vraagt veel creativiteit. Je moet niet vies zijn van heuristieken en simulaties. Als streng wiskundig opgeleide moet je "vuile handen" maken, maar wel op een wetenschappelijk verantwoorde manier. Kun je onderwijzen hoe dat moet? Het lijkt mij lastig. Wel kun je laten zien hoe het moet. En dat heeft Ton zijn hele werkzame leven gedaan. Balancerend op het raakvlak van theorie en praktijk.

In zijn CQM-tijd bekeek Ton aanvankelijk slechts het fabricagestuk van de logistieke keten. Bijvoorbeeld in het project dat hij uitvoerde bij VALVO (Philips) Hamburg. In CQM- Informationsheet December 1988 schrijft hij:

Valvo RHW Hamburg manufactures bipolar integrated circuits. These are applied worldwide in the consumer electronics and automobile industries. The complex production process, a high innovation rate, the many different products and strong competition on the market make it difficult to realize a stable production process. Large stocks of work in progress (WIP) have hitherto been maintained to cope with the inherent uncertainties. Added to the problem of overloaded machines, this has led to long and wildly fluctuating throughput times and -consequently- to poor due date performance.

Firstly, a global capacity tool was developed and installed. It checks whether the available capacity suffices and indicates where extra capacity should be made available. Secondly, it was proposed to install a WIP limit. A simulation study has been carried out to find the WIP limit that leads to better and less fluctuating throughput times.

Het project in Hamburg en ander CQM-werk inspireerden Ton tot wetenschappelijke publicaties. We noemen een artikel in het *International Journal of Production Research*, getiteld "Computationally efficient approximations for balanced flowlines with finite intermediate buffers” (De Kok 1990a). En in het *European Journal of Operations Research* verscheen: “Hierarchical production planning for consumer goods” (De Kok 1990b).

Achteraf gezien kun je dit werk als klein bier beschouwen. Want je wilt vanuit je wiskundige achtergrond vooral een bijdrage leveren aan het optimaliseren van de logistieke keten als geheel. Ton liep daar tegenaan in een project bij Philips Telecom Systems in Hilversum. Samen Henk Zijm publiceerde hij over dat project in *Engineering Costs and Production Economics* (Zijm and De Kok 1989). In de inleiding van het artikel lezen we:

The objective of the study was to carefully review production and inventory control rules and, in particular, to quantify the possible reductions of stock levels as a result of certain well-defined actions, to be initiated by the responsible management.

Hence, we are dealing with a complicated multi-stage, multi-items production process.

Dat was in 1989. In de jaren daarna gaat Ton pas echt los. Een indrukwekkende lijst van publicaties getuigt ervan. Jan van Doremalen zegt over de bijdrage van Ton de Kok:

Kern van Ton zijn bijdrage aan de supply chain wereld vind ik persoonlijk nog steeds zijn gedreven betoog voor een op stochastische modellen gebaseerde blik op de supply chain versus een op deterministische modellen gebaseerde blik. Voor Ton is de supply chain een dynamische zich in de tijd ontwikkelend proces in een onzekere wereld. Het is dan ook zaak om goede beslissingen te nemen met een langere termijn blik, zodat een consistente en goed gebalanceerde keten ontstaat. Voor een goede materiaalstroom-beheersing is het verder essentieel om een multi-echelon bril op te zetten, te streven naar transparantie over de keten heen en beslissingen te nemen die integraal en op langere termijn tot de beste beschikbaarheid en de laagste kosten leiden. Deze visie is natuurlijk in lijn met het bredere gedachtegoed van de Eindhovense school van onder andere Will Bertrand, Hans Wortmann en Jacob Wijngaard.

Collaborative Planning

Als de opeenvolgende schakels in een logistieke keten hun planningen goed op elkaar afstemmen zal het de efficiency van de keten als geheel sterk verhogen. Hoe ondersteun je dat? Rond de eeuwwisseling had Ton de Kok een wiskundig model uitgedokterd om zo’n afstemming van planningen te optimaliseren: het

20-jarig jubileum CQM in het Evoluon. Onder het toeziend oog van Babette de Fluiter en Heleen den Hengst probeert Ton een kortste route te vinden.

collaborative planning model. Philips Semiconductors wilde dat model graag implementeren en CQM werd gevraagd een software-gereedschap te maken gebaseerd op het model van Ton de Kok. Fred Janssen, gepromoveerd bij Ton en daarna in dienst getreden bij CQM, toog aan het werk.

Op enig moment ontstond bij de 3 participanten, TU/e, Philips Semiconductors en CQM, het idee om het gereedschap ook geschikt te maken voor andere bedrijven. Dat zou een flinke investering vergen, waarvan elk der 3 deelnemers 1/3 voor zijn rekening zou nemen. Vervolgens zou de naar verwachting zeer hoge opbrengst in dezelfde verhouding verdeeld worden.

CQM oordeelde dat de hoge investering voor het kleine bedrijf te riskant was. Als de commercialisering van het gereedschap onverhoopt zou mislukken zou CQM failliet gaan en zouden alle medewerkers op straat staan. Philips Semiconductors kocht toen van CQM een flink aandeel in de potentiële, toekomstige opbrengst en betaalde in ruil daarvoor alle uren die CQM in het project stak.

Na enige tijd werd duidelijk dat alleen CQM te zijner tijd het gereedschap in de markt moest zetten en onderhouden. Dat wilde CQM niet. Daarvoor waren de medewerkers niet in de wieg gelegd en aan marketeers ontbrak het al helemaal.

In 2002 werd dan ook een punt achter de commercialisering van het gereedschap gezet. Mathieu Clerkx (Philips Semiconductors) verzuchtte: “CQM is de enige die financieel beter is geworden van dit project”. Dat was waar. Semiconductors en de TU/e verdienden niks **aan** het gereedschap. Maar wel **met** het gereedschap. Philips Semiconductors leverde het tientallen miljoenen euro aan besparingen op. En de TU/e (lees: Ton de Kok) vergaarde wereldwijde roem. In *Interfaces* verscheen het succesverhaal “Philips Electronics synchronizes its supply chain to end the bullwhip effect” en dat resulteerde in de nominatie voor de prestigieuze Franz Edelman Award (De Kok et al. 2005).

Tot slot, Logistics Management Systems (LMS)

Van 2004 tot 2019 werd ik door de TU/e ingehuurd voor de praktijkbegeleiding van LMS-studenten. Zij volgden bij Ton een college “Stochastische processen in de Logistiek”. Met ontzag spraken ze over de geleerde professor. Trots vertelde ik dan dat Ton ooit mijn medewerker was en nu mijn baas. Of LMS wel helemaal onder OPAC viel? De exacte hiërarchische verhoudingen binnen de TU/e bleven mij duister. In ieder geval heb ik nooit last gehad van Ton. Dat had hij geleerd bij CQM. Een goede baas is geen bemoeial.

Eindhoven, november 2024

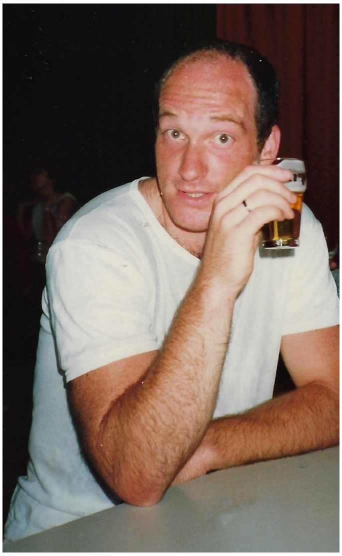

Ton, altijd een innemend mens geweest en gebleven.

Empirical evaluation of some inventory models

Rommert Dekker

1 Introduction

Although Ton always aims for nice mathematical models and analysis, he is also very much concerned about the applicability in practice. In 2024 he pushed me to give a presentation at the ISIR conference in Budapest on my experiences with applying inventory models. This contribution is a write-up of the presentation. We will discuss some standard models, i.e. the EOQ and the Joint Replenishment model. The issues relate very much to the interpretation and estimation of the variables involved.

2 The EOQ model

The Economic Order Quantity (EOQ) model is perhaps the most famous and the most misused model in inventory control. It is an elegant model that balances the costs of having inventories with the burden of a replenishment order. The question we address in this section whether it reflects actual practical situations and whether students and practitioners interpret the model correctly. The assumptions underlying the model are as follows (see also Silver et al. (1998)): (a) demand is constant and known in time, (b) replenishments can be done at any moment in time and are fulfilled in their entirety and instantaneously, (c) the purchase cost of the items is constant in time and in the number ordered (d) the costs of having inventories is linearly related to the capital invested in them and (e) the effort (cost) related to placing and fulfilling a replenishment order does not depend on other items and are consist of a fixed part as well a part which is linear in the number of items ordered (f) no shortages are allowed, (g) the number of items ordered can be any real value (h) we can take a long-run perspective.

As is usually the case, none of these assumptions is likely the case in reality, yet in many cases, the model does its job and the deviations from the assumptions have minor effects. One exception is the case of discounts if more than a certain amount is ordered. We will discuss this later. The EOQ model is simple and serves as a typical mathematical basic step to go from a one-product model to multiple products.

2.1 The formula

The EOQ model is intended to give advice about how many items of one product (SKU) to order in a replenishment. It uses the following quantities for a particular SKU, see table 1. The EOQ is then given by the following formula

Table 1: Nomenclature

D	Demand in items per time unit
S	Fixed ordering costs of a replenishment
h	Holding cost rate per time unit
C	Costs of one item

$$EOQ = \sqrt{\frac{DS}{hC}} \tag{1}$$

The formula has been developed several times in history (see Andriolo et al. (2014) for an overview) and it is contributed to Harris (1913), Wilson (1934) and Camp (1922). While the demand D and costs of one item are easily determined, there are issues concerning the estimation of the fixed ordering costs and the holding cost rate in practice.

2.2 The issues

Fixed ordering costs S

These are the fixed part of the costs for a company to place a replenishment order. It consists of : the work in placing an order (low to zero if orders are placed automatically), the transportation costs, the cost of checking the item, the cost for handling the item (putting it into a warehouse) and paying the order. Both the supplier and the buyer incur costs for these aspects. In international trade the allocation of these costs between these two organizations is arranged in the so-called Incoterms (see Incoterms).

As transportation and handling are often important cost drivers, one usually orders several products at the same time. Hence, ordering costs are split up over many products and it is difficult to assess the fixed order costs for a single product. Moreover, customers do not want to pay explicitly for these fixed costs. Accordingly, the fixed costs are often hidden in the prices of products and discounts are given if many items of the same product are ordered. One can do a time study to estimate the handling time and related costs, yet there are thousands of products and there is quite some variability in the amount of time needed. Students have often taken the total transportation costs as individual fixed order costs with devastating results.

Anecdote when the EOQ model was introduced in 1960-ties in industry, administration offices (the original Dutch name for consultants) had no computer nor calculator but only a slide rule. In one case, they had problems with calculating a square root and simplified it by dividing by 2 as $\sqrt{4} = 2 = 4/2$ and $\sqrt{5} = 2.2 \approx 5/2$. Only after a year did the error became known.

The holding cost rate h

This factor expresses the burden of having inventory in terms of: capital employed, storage space, insurance and obsolescence costs. Standard, it is expressed as a fraction h of the price C of an item. An alternative to this in case storage space is expensive, is to set costs related to the volume of a product. This would apply to very cheap and bulky items, like toilet paper and cotton wool, but we hardly see that in practice and manual overwrites of EOQ quantities occur.

Anyhow, the use of a holding cost rate is an important and very effective tool to manage inventories by models. It is therefore strange that many managers are not aware of the value their organization uses. Specialists should know and sometimes complex formulas are used, like the Weighted Average Costs of Capital (WACC) (see Chopra and Meindl (2016) p. 283). In budget driven organizations, like governments, capital costs are not or less relevant as it is important to completely use a budget. In those cases one is not familiar with holding costs.

We have encountered cases where a spare parts warehouse got the instructions to reduce its capital involved in inventories. To do so, it put the capital burden to the owners of the installations the expensive spare parts were related to. As a consequence the value 0 was entered as the price of the item in the warehouse' data base. This does explode the value of the EOQ though. It is therefore wise to use a cap on the EOQ value and flag those cases where it is larger than say two years of demand.

Another issue with spare parts is that they can stay for a very long time in a warehouse. Using an average-cost view then conflicts with the write-off of the capital involved in the spare part. In fact, one should then use a lower holding cost rate.

Other remarks

Inventories matter less or not in short-haul transportation. Handling costs do matter, especially for cheap items. Boxes, box-wrapping are applied. In those cases one rounds the order quantity off to appropriate values. In fact, using fractional box quantities may lead to picking errors.

Case study: Pharmacy delivery

Joost, an econometric student, did a case study on the logistics of his parents' pharmacy (Göttgens and Dekker (2007)). It was delivered all weekday nights to respond to changes in demand (pharmacies have a delivery obligation). Each day after closure, the pharmacy gave a replenishment order that was delivered by night shipping. The pharmacy applied the double-door principle. The delivery person had access to the first door but not the second one. He/she delivered the pharmaceuticals in delivery boxes in the hallway and took the empty boxes back to the warehouse. The pharmacy staff then had to put the pharmaceuticals in the right cabinet drawers.

Joost was able to retrieve data on demands of the ten most sold products. He also timed the duration of storing of each pharmaceutical (33 seconds) in the cabinet as well as the time needed to check the order for each product (5 seconds). Using these values he was able to apply the EOQ model an advised to lower the order quantities substantially. His mother was happy as it freed up a lot of space in the cabinet.

Later, the case study was given as real-life exercise to logistic MSc students to apply the EOQ model. They had substantial difficulty in estimating the fixed order costs per product, even though the time measurements were given.

Concluding

The EOQ model is a very nice mathematical model, but to apply it, one needs to understand it as well as its assumptions. When teaching the model to students a lot of attention goes to the mathematics of the model, which overshadows the issues with checking its assumptions that come forward when applying it.

3 The Joint Replenishment Problem (JRP)

The joint replenishment model is an extension of the EOQ model to multiple products. There is both a continuous as well as a discrete-time variant of the problem.

3.1 The continuous-time JRP model

Assume we have N products indicated with index $i = 1, ..., N$. Demands, costs are like in the EOQ model and are specific per product. They are indicated with the index i. The holding cost rate is assumed to be the same for all products, as only product costs are relevant. The fixed replenishment order costs are now split-up into a common part S and a product specific part $s_i, i = 1, ..., N$. It is very difficult to determine the overall optimal policy, hence one restricts oneselves to structured policies. In one such policy one has the option to order every T time units at costs S and use that option for product i every k_i'th occasion at costs s_i. If none of the products is ordered one may exclude the major set-up costs S. Combining all, leads to the following formula for the long-term average costs $g(T, k_1, ..., k_N)$

$$g(T, k_1, ..., k_N) = \frac{\Delta(k_1, ..., k_N)S}{T} + \sum_{i=1}^{N} \frac{s_i}{k_i T} + hC_i D_i k_i T \qquad (2)$$

where $\Delta(k_1, ..., k_N)$ indicates the fraction of replenishment opportunities that are being used. If one of the $k_i = 1$, then the delta factor is one, implying that all opportunities are being used. It can be derived from the principle of inclusion and exclusion. It is defined by

$$\Delta(k_1, ..., k_N) = \sum_{i=1}^{N} (-1)^{i+1} \sum_{\{\alpha \subset \{1,...,N\} : |\alpha| = i\}} \left(\mathrm{lcm}(k_{\alpha_1}, ..., k_{\alpha_i})\right)^{-1} \tag{3}$$

$$= \sum_{i=1}^{N} \frac{1}{k_i} + (-1)^1 \sum_{(i,j) \subseteq \{1,...,N\}} \frac{1}{\mathrm{lcm}(k_i k_j)} \tag{4}$$

$$+ (-1)^2 \sum_{(i,j,l) \subseteq \{1,...,N\}} \frac{1}{\mathrm{lcm}(k_i k_j k_l)} \tag{5}$$

$$+ ... + (-1)^{N+1} \frac{1}{\mathrm{lcm}(k_1, ..., k_N)} \tag{6}$$

Several authors have worked on it and in course of time many papers have appeared, presenting both heuristics as well as exact methods with and without the correction factor see Porras and Dekker (2008).

3.2 Issues with the JRP

The same issues for the holding costs for the EOQ apply to the JRP. The same holds for the fixed ordering costs, albeit slightly more complex. The major set-up costs S typically covers the transportation costs as well as cost for billing. The so-called minor set-up costs s_i for product i relate to its handling costs, e.g. putting it at the right place in the warehouse. The latter may be quite difficult to establish: is it a full pallet or mixed pallet load? Do products need to be moved from a pallet by hand? Getting this info by hand is troublesome for a warehouse with more than 10,000 products. Next, even the major set-up costs may be problematic to get as a truck may visit several customers in one round. Moreover, in retail distribution customers want regular distribution and the supplier has to decide on which days to visit which customers. Hence few options exist for the supply interval T.

Additional issue: size of transportation unit

A major requirement is that the whole order has to fit in a single transportation unit. This can be a truck (van, city truck, truck-trailer or megatrailer) or a container (20ft,40 ft or 45 ft). If the total volume (pretty often it is volume and not weight that puts a constraint) exceeds the volume of the transportation unit, either a larger truck or container has to be selected or a second truck. In the latter case costs double. This directly destroys the nice coordination.

Case: Lubricants resupply

A lubricant trading company in Mexico City sourced its lubricants from a refinery in Houston. A large truck was used to pick up the products. These were stored in a variety of packaging, ranging from drums of which four could be put on a pallet, IBCs (which is a plastic container in a metal framework fitting on a pallet) and bottles. Some of these could be stacked, while others not. It will be clear that loading the truck requires a kind of load planning optimisation.

The major set-up costs consist of the transportation costs of the truck going from Mexico to Houston and back. The minor set-up of costs consist of loading and unloading the pallets from the truck into the warehouse. A complicating factor is that one would like to store pallets with one type of drum only, in order to facilitate retrieval. Pallets with several products would need to be split up, increasing the work. It will be clear that the standard JRP would have to be extended with a load planning module, which would make the overall problem very complex.

3.3 Alternative methods

An often used alternative to the JRP model is to set minimum order quantities (MOQ) for individual products, or to use a staircase product cost function. We have encountered a case where the price was increased in 8 stairs. The problem was not to determine the optimum order quantity, as a simple enumeration would do the trick, but how did one ever design such a staircase? Would there be several steps in the fixed order costs? In Porras and Dekker (2006) the JRP problem with MOQ values was analysed. The problem was split-up into two steps. In the first step a common base interval T is considered while in the second step the optimal order quantity given this base interval is determined for each individual product. It is very easy to do the latter step with minimum order quantities. The crux of the approach is to determine when a change of the common base interval changes the individual order frequencies. It appears to be easy to do so, also for minimum order quantities. Given bounds on the common base interval then yields an exact algorithm.

4 Split deliveries and break quantities

In some cases orders are placed in large quantities, for example when a low price is used in the EOQ formula. This may deplete inventories substantially. Ton de Kok worked on two variants of this case. In the first case, the large quantity is split up into several deliveries, hence overcoming the large load on the local inventories. In the second case, the order is transferred to a higher inventory echelon, using a break quantity or cut-off order size to determine when this happens. Both are clever inventory rules used with some companies. The issue with modeling, however, is to find a good cost structure. Two of Ton's papers are about these approaches Janssen et al. (2000) and Dekker et al. (2000).

5 Apology

The EOQ and JRP models are not wrong. Models are never wrong, they are sometimes not appropriate as they may not capture the essence of a problem. These models have been formulated for mathematical beauty and studying them has been a pleasure. Yet, as so often when trying to model actual logistic situations, these situations are highly complex and existing models have to be adapted to get meaningful advice out of them.

6 Ton de Kok

Ton and I both studied applied mathematics in Leiden. Ton started one year after me, but finished shortly after. After our studies we both continued for a Phd. I stayed in Leiden and Ton moved to the VU. We did meet each other in the Phd classes of Prof. Cohen who always suprised us with fully written blackboards. We jointly visited a PhD summer course in Durham, where Ton demonstrated his soccer skills. When our Phd appointments were about to finish there were hardly any positions at universities. It appeared that we both applied for a position at CQM, the OR consultancy group of Philips. As I had military obligations and Ton not, I chose for Shell Research who declared me vital for their operations and Ton got the job at CQM. Seven years later almost the same happened when we wanted to return to academia: both the EUR, the TUE and the UvA had vacancies for a professor in OR. We both ended at our preferred choices and remained friends.

We both worked on inventory control, though I much on single echelon and Ton at multi-echelon. When we met occasionally, I was always triggered by Ton's enthusiasm about his theoretical and practical work and that remained through all his career.

Balance allocations of shortages using QP

Nico Dellaert

Abstract: *In this chapter we consider rather unknown work from the PhD thesis of Judith Spitter (Spitter (2005)). We compare the performance of strategies with an allocation policy with the standard Linear Programming (LP) strategy for capacitated supply chains. The overall results show us that models with allocation strategies perform better than the standard LP model, whereby the model with quadratic objective function performs best. The model with quadratic objective function not only gives lower inventory cost, it also can be solved more efficiently.*

Introduction

It was probably 1985, but it could also have been 1986, when I first met Ton. I had just started my PhD and, as most of my colleagues, I attended the yearly meeting in 'de Blije Werelt' in Lunteren. It was great to see so many passionate OR researchers, both in the interesting talks as well as at the bar and in the forest walks. There were several new guys to me that took my interest and sympathy, especially after the long train trip home, which was delayed severely by the snow. The guy that was able to put the most attention on himself was (probably unsurprising for the reader) of course Ton de Kok. I thought that it could be great fun to work together with such a character. That opportunity came earlier than expected when two CQM employees partly joined the group of Jaap Wessels to discuss production and inventory problems. Next to Ton, we also had Henk Zijm as our inspiration. In 1989, I joined the MABES group of the Econometric Institute of the Erasmus University, at that time led by Alexander

Rinnooy Kan and soon afterwards by Luk Van Wassenhove, and I spent 9 nice years over there, but due to budgetary cuts, the group had to shrink and shrink. By 1998, I saw a job opening at the LBS/OPAC group and, remembering the inspiring figure of Ton, I 'took' this job and became his colleague. That led to very nice collaborations over the years and in the next section, I will pick a few.

Collaborations

Ton was the promotor of a number of 'my' PhDs: Judith Spitter, Joost de Kruijff and Alireza Hesaraki and together with Tom Van Woensel, we co-supervised Duygu Tas and Derya Sever. But we also worked together with other guests or PhDs, like Connie Kohler Gudum, Wang Wei, Fred Janssen and Sjors Jansen. Some of this work got a lot of attention; for instance, the work on stochastic travel times with Duygu and the dynamic shortest path problems with Derya. Also, the work with Alireza gets quite some citations. For other work, the number of citations does not seem to be in line with its quality. I expected that the Safety Stock Adjustment Procedure with Connie would have become a seminal paper, but apparently it was not recognized as such, as it does not even contribute to Ton's h-score. And then, there is the work that got very little attention. Many beautiful ideas did not make it into a published paper, for various reasons. Sometimes, tackling the follow-up challenges proves to be far more demanding, and often more rewarding, than writing and submitting the paper itself. At times, we were disappointed by the limited vision of the referees; other times, different factors played a role. In the next section, I want to pick up a 'forgotten pearl'. It is based on Chapter 6 of the thesis of Judith Spitter(Spitter (2005)).

Balance allocation of shortages

In SCOP planning, we often notice an imbalance in the allocation of materials. This imbalance is also detected in De Kok and Fransoo (2003). They also showed that the LP-based SCOP function is outperformed by a so-called synchronized base stock (SBS) policy developed by De Kok and Visschers (1999). The SBS policy uses allocation mechanisms derived from the analysis of divergent systems and general supply chain structures. However, the results of De Kok and Fransoo (2003) are restricted to uncapacitated supply chains because the SBS policies do not take into account capacity restrictions. The focus of this chapter is to develop LP models for capacitated SCOP problems with allocation strategies. We restrict ourselves to material allocation and assume a one-to-one relation between item and resource. The introduction of allocation strategies into the LP formulation implies that we add additional constraints. Therefore, the optimal solution to the new problem may have a higher cost than in the original formulation. However, the LP model is part of a rolling schedule implementation. We will show that, despite the increase in costs on an instance, the actual costs incurred under uncertainty can be substantially lower.

Allocation problems occur only in the divergent parts of arbitrary supply chains. For pure divergent systems, allocation rules have received considerable attention. Eppen and Schrage (1981) introduced a fair share allocation rule, based on equal stock-out probability at end-stock points, for a two-tier system without intermediate stocks. In De Kok (1990a) the Consistent Appropriate Share (CAS) rationing policy is introduced, with the fill rate (the fraction of demand delivered immediately from the stock on hand) as a service criterion for 2-echelon systems. A generalization of the CAS rationing policy for two-echelon divergent systems is the Balanced Stock (BS) rationing introduced by van der Heijden (1997). In van der Heijden et al. (1997) this allocation rule is extended for general N-echelon distribution systems where all upstream, downstream, and intermediate stock points are allowed to hold stock. In De Kok and Fransoo (2003), a correction of the formula for allocation fractions is given for non-identical successors.

In this chapter we extend the SCOP model with linear allocation rules. These linear allocation rules divide shortages among the parent items by predefined allocation fractions. Furthermore we introduce a model, whereby the linear objective function is replaced by a quadratic objective function. The idea behind this is, that similar to the linear allocation rules, the quadratic objective function ensures a balanced rationing of the shortages among the parent items. Under a linear objective function shortages are typically rationed among "cheap" parent items, while "expensive" items do not take part in the shortages at all.

In De Kok and Fransoo (2003), it is shown that the original SCOP model is outperformed by the so-called SBS policy, developed in De Kok and Visschers (1999). In this chapter, we compare the SBS policy with the LP models introduced with allocation strategies. We see that the SBS policy still performs a little better, but the differences are small. Since the SBS policy is only suitable for uncapacitated supply chain structures, the introduced LP models with allocation strategies provide the generality needed to handle arbitrary capacitated supply chains. Hence, at the end of this chapter, we also measure the performance of the original LP model with the LP models with allocation strategies for capacitated supply chains and various allocation fractions.

Extended model

By adding constraints to the problem, the LP compensates for any shortage with imaginary slack, so that the order requests are not restricted. So, we have to penalize slack by adding a cost component for slack in the objective function. The penalty for imaginary slack must be high enough so that it is not beneficial to create it, but must not be so high that having slack is more costly than having backorders. Thus, the related costs should be smaller than the minimum backorder costs, but cannot be too small. Hence, we subtract small ε costs from the backorder costs and use that as penalty costs for creating imaginary slack. So, we consider the following objective function

$$\min \sum_{t=1}^{T} \sum_{i=1}^{n} \alpha_i I_{i,t-1} + \sum_{t=1}^{T} \sum_{i=1}^{n} \beta_i B_{i,t-1} + \sum_{t=1}^{T} \sum_{i=1}^{n} \left(\min_{j \in E_i} \beta_{jt} - \varepsilon \right) \left(S_{it}^{+} + S_{it}^{-} \right), \quad (1)$$

where ε is a very small number.

Allocation fractions

The determination of allocation fractions, q_{ij}, has received considerable attention in the literature. We consider four different allocation fractions. The first allocation fraction is based on the number of parent items; shortages are allocated proportionally to the number of parent items by

$$q_{ij}^{|N|} = \frac{1}{|\operatorname{succ}(i)|} \quad i = 1, \ldots, n, j = 1, \ldots, n \quad (2)$$

In the second allocation fraction, we also take the mean demand of the parent items into account. This allocation fraction with

$$\mu_j = \sum_{k \in E_j} E\left[D_k\right] \quad j = 1, \ldots, n$$

is given by

$$q_{ij}^{\mu} = \frac{\mu_j}{\sum_{k \in \operatorname{succ}(i)} \mu_k} \quad i = 1, \ldots, n, j = 1, \ldots, n \quad (3)$$

The third allocation fraction arises from Eppen and Schrage (1981). They introduce a fair share allocation rule for a two-echelon system without intermediate stocks. The allocation rule ensures that the end-stockpoint probabilities are equalized. Defining

$$\sigma_j = \sigma\left(\sum_{k\in E_j} D_k\right) \qquad j = 1,\ldots,n$$

the allocation fraction of Eppen and Schrage (1981) is represented by

$$q_{ij}^{ES} = \frac{\sigma_j}{\sum_{k\in \text{succ}(i)} \sigma_k} \qquad i = 1,\ldots,n, j = 1,\ldots,n. \tag{4}$$

van der Heijden (1997) determined linear allocation policies that minimize the probability of imbalance. Imbalance is the allocation of a negative quantity from an intermediate stockpoint to at least one of its successors. De Kok and Fransoo (2003) corrected an error in the analysis of van der Heijden (1997) and obtained the following expression for the allocation fraction.

$$q_{ij}^{HKF} = \frac{\sigma_j^2}{2\sum_{k\in \text{succ}(i)} \sigma_k^2} + \frac{\mu_j^2}{2\sum_{k\in \text{succ}(i)} \mu_k^2} \qquad i = 1,\ldots,n, j = 1 \tag{5}$$

Quadratic objective function

In the previous sections, we balanced the allocation of shortages by adding allocation rules to the LP model. In this section, we balance the allocation by replacing the linear objective function with a quadratic objective function.

The objective function becomes

$$\min \sum_{s=t+1}^{t+T}\sum_{i=1}^{n} \tilde{\alpha}_i I_{i,s-1}^2 + \sum_{s=t+1}^{t+T}\sum_{i=1}^{n} \tilde{\beta}_i B_{i,s-1}^2, \tag{6}$$

where $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ are constants to indicate the relations between inventory and backorder costs of the items.

Weight factors in quadratic objective function

In this section, we consider three different kinds of weight factors. The first set of weight factors is equal to the cost structures used. Thus,

$$\begin{aligned} \tilde{\alpha}_i &= \alpha_i \quad i = 1, \ldots, n \\ \tilde{\beta}_i &= \beta_i \quad i = 1, \ldots, n. \end{aligned} \tag{7}$$

The other two sets of weight factors are inspired by the allocation fractions of Eppen and Schrage (1981) and of van der Heijden (1997) and De Kok and Fransoo (2003), so instead of only cost, we also take the demand distribution into account. The weight factors are chosen such that the inventory and the backorder costs of one item per period are higher for items with smaller allocation fractions. The weight factors are given by

$$\begin{aligned} \tilde{\alpha}_i^{ES} &= \left(\min_{j \in \text{ prec } (i)} \frac{\sum_{k \in E_j, k \neq j} \sigma_k}{\sum_{k \in E_j} \sigma_k} \right) \alpha_i^2 \qquad & i = 1, \ldots, n \\ \tilde{\beta}_i^{ES} &= 10 \tilde{\alpha}_i^{ES} & i = 1, \ldots, n. \end{aligned} \tag{8}$$

and

$$\begin{aligned} \tilde{\alpha}_i^{HKF} &= \left(\min_{j \in \text{ prec } (i)} \frac{\sum_{k \in E_j, k \neq j} \sigma_k^2}{\sum_{k \in E_j} \sigma_k^2} + \frac{\sum_{k \in E_j, k \neq j} \mu_k^2}{\sum_{k \in E_j} \mu_k^2} \right) \alpha_i^2 \\ \tilde{\beta}_i^{HKF} &= 10 \tilde{\alpha}_i^{HKF} \qquad & i = 1, \ldots, n. \end{aligned} \tag{9}$$

MP models with balanced allocation vs SBS concept

De Kok and Fransoo (2003) show that synchronized base stock (SBS) concepts perform better than existing LP models for uncapacitated supply chains with stochastic demand. In this section, we compare this SBS concept with the mathematical programming models introduced with allocation strategies. In the first section, the SBS concept is discussed, and in the next section, the experimental setting, also used in De Kok and Fransoo (2003), for comparing the model is given. We conclude with the results of the comparison.

Synchronized base stock control concept

The synchronized base stock (SBS) control concept, introduced by De Kok and Visschers (1999), is a class of policies that can be applied to general supply chain structures. The allocation mechanisms used in the SBS concept are derived from the analysis of divergent systems (cf. De Kok and Fransoo 2003). With the SBS concept, it is possible to characterize the optimal policy for non-capacitated situations under i.i.d. exogenous demand, and near-optimal policies can be found numerically.

From De Kok and Fransoo (2003) we learn that the SBS concept considerably outperforms the LP-based concept. An explanation of the superiority of the SBS concept was found in the way LP tends to prioritize items in case of a shortage of upstream availability rather than rationing among the items that need this upstream availability. The models proposed in this chapter have allocation strategies for rationing these shortages. Hence, we compare the stochastic SBS concept with the deterministic MP models with allocation strategies using a rolling horizon.

In the numerical comparison, we restrict ourselves to the situations with infinite resource availability, since the SBS policy cannot deal with finite resources. The results of capacitated systems are only available for single item, single stage systems (see e.g. De Kok 1989), serial systems (Tayur 1993), or for divergent systems where only the most upstream stage is capacitated (see De Kok 2000).

We will first consider a general system with 11 items. For the planned lead times, we have, analogously to the cost structure, the following variables:
$\tau_i = \tau_f$ planned lead time end-items, i $= 1, 2, 3, 4$
$\tau_i = \tau_s$ planned lead time specific components, i $= 5, 6, 7, 8$
$\tau_i = \tau_{sc}$ planned lead time semi-common components, i $= 9, 10$
$\tau_i = \tau_c$ planned lead time common components, i $= 11$

We vary the planned lead times $(\tau_s, \tau_{sc}, \tau_c)$ as follows $(1, 2, 4), (4, 2, 1)$ and $(1, 4, 2)$. The safety stocks are chosen such that we obtain a non-stockout probability of 95%.

Results

The results of the comparison between the SBS concept and the LP models with allocation strategies are given in Table 1. Although the results of the standard LP model are better than in De Kok and Fransoo (2003), the difference between the SBS concept and the standard LP model is still significant, and the use of allocation strategies could give better results.

		Supply chain inventory cost						
cv_i^2	$(\tau_s, \tau_{sc}, \tau_c)$	SBS	LP_{st}	Δ_{LP}	LP_{alloc}	Δ_{alloc}	QP	Δ_{QP}
0.25	$(1, 2, 4)$	71682	79477	10.9%	73458	2.5%	73332	2.3%
0.25	$(4, 2, 1)$	76476	78133	2.2%	78853	3.1%	76765	0.4%
0.25	$(1, 4, 2)$	73550	80620	9.6%	74197	0.9%	74172	0.8%
0.5	$(1, 2, 4)$	104448	115227	10.3%	106923	2.4%	106679	2.1%
0.5	$(4, 2, 1)$	112316	114659	2.1%	115696	3.0%	113381	0.9%
0.5	$(1, 4, 2)$	107616	115386	7.2%	108376	0.7%	108362	0.7%
1	$(1, 2, 4)$	152203	168533	10.7%	158101	3.9%	157752	3.6%
1	$(4, 2, 1)$	165328	169122	2.3%	170590	3.2%	167180	1.1%
1	$(1, 4, 2)$	157034	169030	7.6%	159588	1.6%	159722	1.7%
2	$(1, 2, 4)$	218551	247578	13.3%	233044	6.6%	232392	6.3%
2	$(4, 2, 1)$	245998	249849	1.6%	250075	1.7%	247705	0.7%
2	$(1, 4, 2)$	228789	247571	8.2%	235761	3.0%	235694	3.0%

Table 1: Inventory costs of different SCOP functions.

Looking at the results, we see that with uncapacitated resources, the allocation strategy with quadratic objective function outperforms the strategy with allocation rules. Furthermore, we see that for the cases where specific items have a planned lead time of four periods, the performance of the quadratic model is only a little better, and the performance of LP_{alloc} even worse than the performance of the standard LP model. If the specific items have the longest planned lead time, they are the bottlenecks of the system; this means that the (semi-)common items do not have an allocation problem. So, for supply chain structures where shortages can be allocated to several parent items, linear allocation rules improve performance. But in those supply chain structures without an actual allocation problem, interference of allocation rules diminishes the performance of the LP model. On the other hand, for problems where the common item has the longest planned lead time, the common item is the bottleneck, and allocation strategies improve the performance of the LP model.
In the last column of Table 1 , the relative difference between SBS_{sim} and the best allocation strategy, namely QP, is given. We see that SBS policies still outperform mathematical programming models, but the differences are small. The largest differences are found when the common planned lead time is equal to four periods, so apparently the SBS concept better exploits the commonality. From the discussion in De Kok and Fransoo (2003) it follows that this may be caused by better parameter settings. The echelon stock control policy allows for implementing "push" policies where upstream stocks are zero, while the LP model tends to

build up stock upstream. A possible explanation for this is that the LP model aims at satisfying the end item demand forecast. After satisfying this forecast, the remaining stock of components is not used for the assembly of end items because the stock of components is cheaper than the stock of assembled end items.

Numerical experiment with capacitated supply chain structures

In the previous section, it turned out that the SBS policy still outperforms MP-based concepts. However, the differences are small, and the SBS policy is restricted to uncapacitated problems. Hence, MP-based concepts are good alternatives for capacitated problems. In this section, the performance of the allocation strategies is measured by comparing them with the standard LP model.

Experimental design

Allocation problems occur only in divergent parts of the supply chain. Hence, for the experiments, we need a supply chain with a divergent part. We have chosen a so-called "W-structure," whereby we have two end items, and each of these end items consists of a common and a specific item. We use the "W-structure" instead of a pure divergent structure to show that the allocation strategies also work if the allocation of common child items is restricted by the number of specific child items.

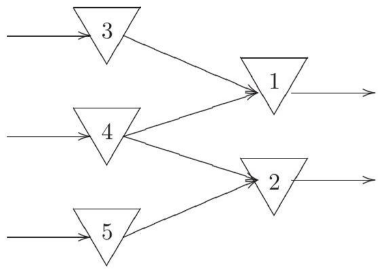

Figure 1: Schematic representation of a 2-echelon system.

We consider a one-to-one relation between items and resources, i.e. each item is produced on its own resource. For simplicity, we number the resources according to the number of the item it produces. The end items are produced on uncapacitated resources, while the specific and common child items are not. In the results of the comparison between the SBS and the MP models with allocation strategies, we have seen that if the specific items have long planned lead times, the allocation strategies did not have any effect on the inventory cost. Long lead times are comparable with high utilization rates of resources. Hence, we restrict the experiments to test cases whereby the common item is the resource with the highest utilization rate. We consider two cost structures and expect large reductions in inventory cost, especially when the costs of end items are equal. To identify more situations in which allocation strategies reduce the overall inventory costs, we also look at various settings.

Results

The performance of each model with allocation strategies is determined by the relative difference in inventory costs between the model with and the model without allocation strategies. In Table 2, the relative average, maximum and minimum differences are given. Note that "QUANT" corresponds with the allocation fraction $q_{ij}^{|N|}$ and "MEAN" with the allocation fraction q_{ij}^{μ}.
If we look at the overall results we see average reductions in inventory costs of more than 20%. The best performance is reached by using a quadratic objective function; the average reduction exceeds 25%. Looking at the maximum reductions, savings of 65% can be made. Unfortunately, there are also cases in which inventory costs become greater when allocation strategies are used. However, the number of cases with negative results is small. We noticed that negative results occur only in cases with totally unequal parameter settings of the end items and high utilization rates of the resources, whereby the highest costs correspond with the item having the biggest demand and highest variation. Apparently, too many shortages are allocated to these items when using allocation strategies. Using a quadratic objective function with weight factors based on the allocation fractions discussed in Eppen and Schrage (1981) and in De Kok and Fransoo (2003), no negative results occur at all.

Furthermore, we see that allocation fractions that also take into account the variance in demand perform better than both $q_{ij}^{|N|}$ and q_{ij}^{μ}. Especially in the worst-case scenario with unequal demand and coefficient of variance, the results are depressing.

	Average	Maximum	Minimum
Allocation QUANT	18.83%	65.43%	−48.65%
Allocation MEAN	20.27%	65.51%	−26.43%
Allocation ES	21.66%	65.45%	−4.73%
Allocation HKF	21.79%	65.51%	−14.85%
Quadratic	24.28%	66.94%	−13.49%
Quadratic ES	26.73%	67.60%	0.71%
Quadratic HKF	26.98%	67.42%	0.68%

Table 2: Average, maximum, and minimum relative difference in inventory costs, compared to the standard LP model.

so it is worthwhile to use more sophisticated allocation fractions. The differences between the allocation fractions of Eppen and Schrage (1981) and De Kok and Fransoo (2003) are very small, but in general q_{ij}^{HKF} gives better results. The performance of the lower bound allocation strategies is equivalent to the performance of the model with linear allocation rules. Hence, in the remainder of the chapter, we focus only on the allocation strategy with linear allocation rules.

In models with quadratic objective functions, we distinguished three different sets of weight factors. We see that the weight factors for which we also take the mean and variance of the actual demand into account perform better, but the differences with the first set of weight factors are small (around 2%). Hence, the quadratic objective function is a good and robust alternative to the LP model.

Since we noticed that allocation problems typically occur in situations with equal inventory costs for end items, we separately show, in Figure 2, the performance of models with allocation strategies for equal and unequal inventory costs. Indeed, we see large reductions when the holding costs of the end items are equal, but for unequal holding costs the reductions are still considerable. Notice that the difference between the model with linear allocation rules and the model with a quadratic objective function is greater when the holding costs are unequal. So, when using the linear allocation rules, too many shortages are allocated to the expensive parent item.

In Figure 3, we give the performance of the allocation strategies for different utilization rates. We see that cost savings are made primarily if the utilization rates of the resources are high. Since shortages occur more often when utilization rates are high, savings are greater when allocation strategies are used. Hence, for $(\rho_3, \rho_4, \rho_5) = (0.8, 0.85, 0.8)$ the savings are only 5%, while for $(\rho_3, \rho_4, \rho_5) = (0.9, 0.99, 0.9)$ the savings are more than 35%.

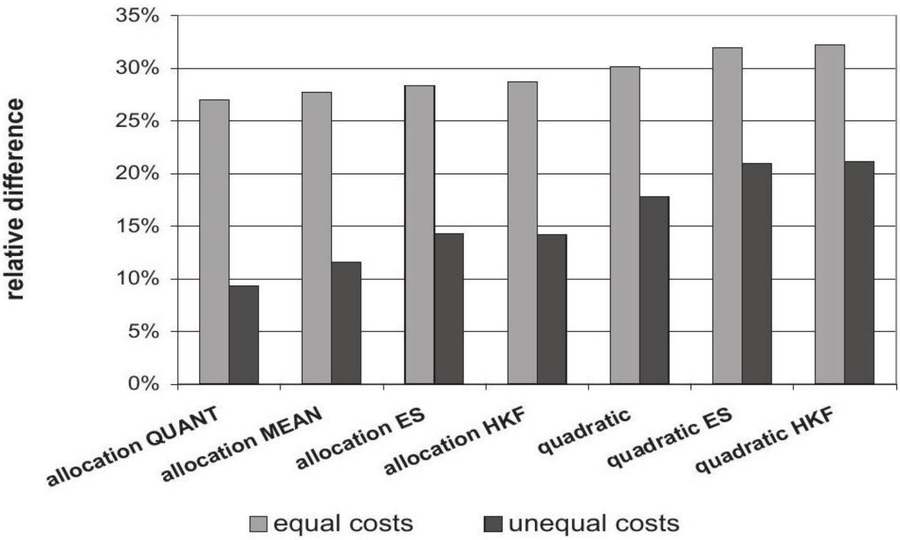

Figure 2: Relative difference between allocation strategies and the standard LP model for equal and unequal costs.

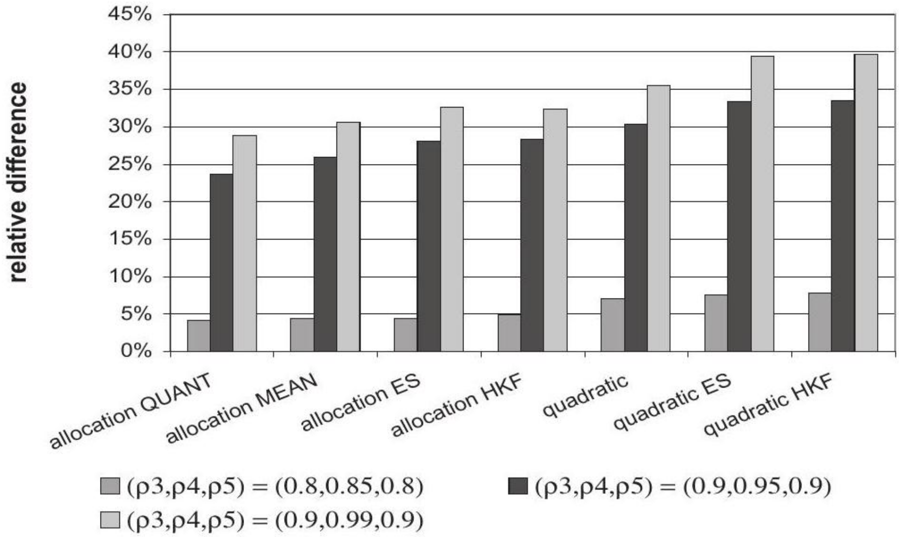

Figure 3: Relative difference between allocation strategies and the standard LP model for different capacities.

We also looked at the time required per test case. Differences between CPU times are less than 25 percent, so not of real importance.

Conclusion

Supply Chain Operations Planning problems of multi-item, multi-echelon systems with demand uncertainty can be solved using Linear Programming models in a rolling-horizon setting. Negenman (2000) and De Kok and Fransoo (2003) noticed that the LP models do not seem to find the right balance in the allocation of stock among the different items. In this chapter, we include material allocation strategies in the LP model by adding linear allocation rules to the LP model. Linear allocation rules allocate shortages of a child item to all its parent items. To identify shortages, the inventory position of the hierarchy is added to the model. In the second model, the linear allocation rules are relaxed by introducing lower bounds for the amount of shortages allocated to the parent items. A third strategy for solving the allocation imbalance is by replacing the linear objective function with a quadratic one. Note that, by the introduction of allocation strategies, additional constraints are added to the LP model. Therefore, the optimal solution to the problem may have a higher cost than the original solution. However, the LP model is part of a rolling schedule implementation. We have shown that, despite the increase in cost on an instance, the actual costs incurred under uncertainty can be substantially lower.

In De Kok and Fransoo (2003) it is shown that for supply chains with infinite capacity and identical end items, the so-called synchronized base stock (SBS) policies outperform MP models. De Kok and Fransoo (2003) state that the concepts based on LP do not seem to find the right balance in the allocation of stocks among the different items. Hence, in this chapter we compared the SBS policy with the introduced MP models with allocation strategies. The results are promising. In particular, the model with the quadratic objective function closes the gap between SBS and LP. Since the model with the quadratic objective function (QP model) is, in contrast to the SBS policy, applicable to general supply chain structures, the QP model is a good alternative for capacitated structures which cannot be solved by the SBS policy.

Furthermore, we measured the performance, compared to the standard LP model, of the allocation strategies for capacitated supply chain structures. For the experiments, we used a so-called "W-structure," with two end items whereby each item consists of a specific and a common child item. The overall results show us that models with allocation strategies perform better than the standard LP model, whereby the quadratic programming (QP) model performs best. The QP model gives not only lower inventory cost, it also is faster in solving the SCOP problem.

Some thoughts on SBS and conventional Production Control Frameworks

Will Bertrand

1 Introduction

Apart from his many managerial and consultancy achievements, Ton de Kok also made impressive contributions to finding a solution to the control of multi-item convergent-divergent production networks under stochastic demand. After detecting several fundamental flaws in the widely used MRP-1 logic for controlling such networks, using his creativity and mathematics skills, Ton in the 1990s finally developed the Synchronized Base Stock system. In tests, the SBS system has been shown to outperform the widely used MRP-1 logic. Comparing the performance of SBS with the performance of the MRP-1 logic for the control of multi-item divergent-convergent production networks, Ton could demonstrate not only that the SBS system produced superior performance, as was his prediction based on the flaws he detected in the MRP-1 logic, but also that the performance improvements can be substantial. This would make the SBS system a natural successor to the MRP-1 logic. From an operations control perspective it is therefore surprising that the MRP-1 logic has not been widely substituted by the SBS system. This may be due to the intuitive appeal of the MRP-1 logic, and its the wide acceptance in industry as part of the MRP-2 control framework. Despite frequent criticism from operations research scholars in the past, MRP-1 and MRP-2 by now are the standard tools for production and supply chain control in industry. MRP-2 evolved out of MRP-1 and as such, these are smoothly integrated systems. This is a substantial advantage in industrial applications. In this note, we therefore investigate if and how the SBS system can be smoothly integrated in the MRP-2 control framework. The note is organized as follows.

In Section 2 we set the stage by discussing the relevant control domains, using the MRP-2 control framework as context. In Section 3 we discuss MRP-1 as the standard materials coordination tool within the MRP-2 framework. In Section 4 we discuss SBS as a materials coordination system within the MRP-2 framework. Section 5 discusses the possible dynamics in the Master Production Schedule to which the SBS system should respond. Section 6 presents the adaptations needed at the organizational and control level to enable embedding of the SBS system in the MRP-2 control framework. Section 7 gives the conclusions.

2 Production Control Domains

Production control systems consist of a hierarchy of related decision functions, each of which makes use of information from separate sources, among which the output of other decision functions, and exercises control over different variables. Together these decision functions constitute the overall control system. One of these decision functions is the Coordination of Production and Sales, CPS for short. The goal of this function is to manage and coordinate the materials control process and the end-product sales process, given the medium-term constraints and flexibilities in the production and sales domains. The flexibility in the production domain is used to adapt the production output to prospective demand and demand constraints, and the flexibility in the sales domain is used to influence demand in line with prospective product output constraints. The actual realization of production output and sales is left to Materials Control, MC for short and Sales Control, SC for short. CPS just coordinates their efforts. The economic perspectives of CPS are return on investment (in materials and capacity) and revenue management. CPS continuously monitors the developments in the market, the competitor's moves, the overall economy, the own sales, and based on these observations determines an optimal response in terms of production output targets and sales targets. This implies that periodically output targets and sales targets can change. The targets decided by CPS are communicated to the production and sales domains.

The most widespread used production control framework in industry is the MRP-2 framework, of which the MRP-1 logic is a part. The MRP-2 conceptual framework can be 'found' in most commercially available Enterprise Resource Planning software packages. The term 'found' refers to the fact that these software packages feature standardized tables and relationships between data in tables that materialize the different MRP-2 concepts. The MRP-2 framework supports a specific way of coordinating production and sales. MRP-2 however does not provide decision support for how to decide about output targets and sales targets. Instead, it supports the CPS function only via the possibility to project the resulting decisions in a table called the Master Production Schedule. The MPS is a Target Production Output statement, which at the same time is a Target Sales statement, over a horizon that covers the stacked lead time of the production system. The MPS procedure itself is not formalized and may consist of a heuristic search for a target output that is not in conflict with materials availability, capacity constraints, and market constraints, and that maximizes future revenues from existing and future production output. We can think of the MPS procedure as a kind of trial-and-error process that can be supported by tools. However, these tools are not generic and are specific for the production/sales situation at hand.

Once the MPS is set, it supports coordination of production and sales in the sense that the production system is assumed to transform the existing materials state into one that is in line with the MPS, and the sales system is assumed to sell the projected production output in the MPS at the best possible price. Using one single MPS for instructing both the production and the sales control subfunction ensures a kind of coordination. The actual actions in the materials control system and in de sales system are not taken by CPS but are delegated to Materials Control and Sales Control. Thus, the relationships between CPS on the one hand, and Materials Control and Sales Control at the other hand, is a kind of 'master-slave' relationship where CPS specifies targets that Materials Control and Sales Control must realize.

The function 'Coordination of Production and Sales' enables a firm to balance its production and sales policies. It is a crucial control function. In most firms that have implemented MRP-2 (or ERP) enterprise software, the coordination of production and sales is achieved by periodically (weekly or monthly) specifying a Master Production Schedule. In this note, assuming that coordination of production and sales is done via periodically specifying an MPS, we discuss if and how the SBS system can be used for materials coordination, or whether the MPS function, or the SBS system, must be changed for the SBS system to be applicable.

3 Materials Control in the MRP-2 framework

How is materials control supported in the MRP-2 framework? In the MRP-2 framework, CPS sets an MPS, being a time-phased target output for each end-product and a Rough-Cut Capacity Plan, which is the time-phased projected capacity load for critical resources. CPS checks for availability of critical resources and makes sure that the MPS is not in conflict with critical resource availability. Non-critical resources are resources that can be easily adapted on short term and are under control of Materials Control. Using its discretion with respect to non-critical resources and materials, MC uses the output of the Materials Requirements Planning calculations (referred to as MRP-1) to adapt non-critical resources and release or reschedule production orders and purchasing orders to move the flow of materials as close as possible the time-phased output target in the MPS.

The MRP-1 logic operates on the actual materials state in the production system and considers batch sizes, safety stocks, and production order throughput times via production lead times. MRP-1 generates actions to be taken to bring the state of the materials in line with the target in the MPS, irrespective of whether the actions are implementable or not. In the MRP-1 logic, orders already released to the shop floor and to the suppliers can be rescheduled, and orders can be released to the shop floor for which materials are not (yet) available. Rescheduling orders are signaled on the so-called exception messages list, and Materials Control is assumed to take actions on them. The MRP-1 logic assumes that flexibility exists at the shop floor level to accommodate these rescheduling actions. In many production systems, some flexibility to reschedule indeed will exist, so the MRP-1 assumptions about shop floor rescheduling capabilities will to some extent be justified. However, shop floor and supplier flexibility will be limited, since it is costly, and not all proposed rescheduling actions will be completely followed up. As a result, if the MPS is dynamic, that is, target data about a specific future period change each time a new MPS is made, MC will not always be able to follow up and the actual periodic product output will deviate from the periodic target output in the MPS.

Sales Control aims at selling the target output in the MPS and thus shortages may occur if output is not equal to target. However, shortages may also occur if output is equal to target and sales is higher than target. Therefore, in the MRP-2 control framework, when using MRP-1 for materials coordination, safety stock of end-products is needed to control delivery reliability. These end-product safety stock norms are inserted in the MRP system before the MRP-1 calculations are made. From literature it is not clear which agent sets these norms. Is it CPS, is it MC or is it SC? Given that CPS coordinates production and sales, and the end-product stock serves as a buffer for both the deviations from the production output targets and the deviations from the sales targets, it is reasonable to assume that, in the MRP-2 control framework, CPS sets the safety stock norms for the end-products.

From the above discussion it follows that in the MRP-2 control framework, Materials Control actions are periodically decoupled from Sales Control actions and vice versa; each agent acts independently based on the same MPS as target set by CPS, the one to realize the output in the MPS, the other to sell the output in the MPS. Moreover, we postulate that in the MRP-2 control framework, CPS sets stock norms for end-products.

4 SBS for materials control in the MRP-2 control framework

Can the SBS system be used to perform the materials coordination function in the MRP-2 control framework? The SBS system aims at achieving a service target for end-products by controlling the base stock of end-products in the presence of stationary stochastic demand for these end-products, while minimizing system-wide stock costs. The SBS system has similarities to MRP-1 in that it uses Bill-of-Materials relationships, fixed lead times for production orders, and acts on the state of the system in terms of stocks and work-in-process. However, SBS does not assume flexibility to exist in the production system. It does not try to speed up or delay already released production orders and only releases new production orders if materials are available. Moreover, SBS does not aim at a target safety stock and a target output for end-products but controls base stocks to realize a target service level for end-products assuming stationary stochastic demand. These two properties are at odds with the MRP-2 control hierarchy in which CPS controls the end-product stocks and specifies each period a target output over the stacked lead time of the production system which Materials Control should realize.

Can MRP-2 and SBS be reconciled by some adaptions in the one or in the other, or in both? Specifically, what setting does CPS create for SBS?

We first look at capacity. With respect to supply, each period CPS checks for availability of critical resources in the Rough-Cut Capacity Plan and sets an end-product output target (the Master Production Schedule) which both are realistic and consistent. Realistic means that both capacity and materials output can be realized with the means available at the production side. Consistent means that the planned capacity is sufficient for the output in the MPS to be realized. We focus on the materials coordination aspect of Materials Control and assume that CPS will always set an MPS for which sufficient critical resource capacity will be available. Thus, in materials coordination it can be assumed that the distribution function of the throughput times of released production orders will be constant over time. In fact, the requirement that the production order throughput time distribution function is constant, is a concretization of the requirement that capacity target and the output target should be consistent. Under the assumption of constant production order throughput time distributions, the constant lead time assumption underlying the SBS system is justified. The SBS system assumes realistic and consistent capacity targets and output targets, and this is exactly what is provided by CPS. In other words, CPS creates a condition necessary for the SBS system to be valid.

Next, we look at materials. CPS each period specifies an MPS, being the output target for Materials Control, which at the same time is the sales target for Sales Control. Assuming that, given appropriate sales effort, the sales target is realistic, it is reasonable to expect that the actual sales will on average be equal to the sales target. However, given the uncertain nature of the market, the actual sales per period will not always be equal to the sales target; there will be some stochastic variation around the target. Moreover, the MPS can vary each period and can differ from the previous MPS. Observations of these variations over time can provide Materials Control with an estimate of the distribution of periodic change in MPS and the deviations of actual sales from the MPS. To Materials Control therefore, the MPS acts as a periodic forecast of future demand.

Thus, if materials coordination would be executed with the SBS system, Materials Control should take the data in the Master Production Schedule as its 'demand forecast'. What does this demand forecast look like? We consider one end-product. Each period t, CPS sets a new MPS for this end-product consisting of a vector $u(t,\xi)$, $\xi = 0,\ldots,L-1$, with L being the stacked lead time of the production system. The number $u(t,0)$ is the output target (and sales target) for period t, determined at the start of period t. The numbers $u(t,\xi)$, $\xi = 1,\ldots,L-1$ could be considered by Materials Control as the current forecast of future demand, knowing that next period there will be a new current forecast of future demand. If the MPS is 'flat', that is: $u(t,\xi) = e(t)$ for all ξ, $e(t)$ can be considered as the expected value of all future demands per period, as formulated at the start of period t. The fact that $e(t)$ varies from period to period implies that Materials Control must take into account the existence of future changes in the level of demand, when controlling the base stocks to achieve a specific service level; just like Materials Control takes into account the existence of future deviations of demand per period from the expected demand per period

Note that the MRP-1 logic assumes that there exists flexibility in the production system that enables MC to reschedule already released production orders and delay or speed up the flow of orders already released. This may not always be realistic. SBS does not assume such flexibility to exist. SBS assumes that production orders are always delivered after the production lead time and acts on the actual state of the materials and generates production orders that can be readily implemented. This is very different from MRP-1 where the actual decisions require an additional layer of human decision making. This extra layer of human decision making is costly in itself and these costs add up to the shop floor rescheduling costs. Note that SBS does not require human decision making or rescheduling.

When using SBS for materials coordination in the MRP-2 control framework, two stochastic variables must be considered. First, the stochastic variation in the level of the demand as given by the MPS, second, the stochastic variation of the demand per period around the target sales level. We next have a closer look at the dynamics in the MPS.

5 Characterizing MPS dynamics

Adapting the MPS to new market information can be done with different agility. In this note we assume no market seasonality or trends. Looking at the current state of the market, and neglecting the current state of the production system, CPS might want to specify a **market optimal target sales level** $s(t)$. This is the sales level that optimizes expected revenues minus sales effort costs. However, this optimal sales level, which simultaneously would be the production system output level, may require a change in output level relative to the output level specified in the previous period, and changing output levels each period might be costly. Thus, considering the change costs, CPS could decide to implement the optimal sales level with some delay of ρ periods. Note however, that this implies that the target sales that CPS specifies for Sales Control in the MPS, are then not optimal i.e., not equal to the market-optimal level $s(t)$ for the pertaining period. Thus, avoiding change costs at the production side implies suboptimal results at the market side.

We can characterize an MPS policy as **conservative** if the delay $\rho = L - 1$, resulting in $u(t, \xi) = u(t-1, \xi+1)$ for $\xi = 0, \dots, L-2$, with $u(t, L-1)$ free to choose by CPS. With a conservative MPS policy, CPS respects the current state of the work-in-process in the production system, completely abstains from disrupting the regular flow and responds to new market information by taking sales measures to sell the products that are in the pipeline. The MPS reflects the optimal sales levels in past periods, leading to: $u(t, \xi) = s(t - (L-1) + \xi)$.

If the sales efforts are successful, the expected demand per period will be equal to $u(t, 0)$ and the safety stock for end-products must be sufficient to cope with a demand uncertainty with average 0 and variance $V(d, u)$. Since the production system stacked lead time is L, the safety stock must be sufficient to cope with a variance equal to $L \cdot V(d, u)$.

We can characterize an MPS policy as **responsive** if the delay $\rho = 0$, resulting in $u(t, \xi) = s(t)$ for all ξ. A responsive MPS policy enables CPS to respond to new market information requiring MC to keep the production system in a state in which it can immediately follow the changes in optimal sales level $s(t)$. Since SBS does not assume production system flexibility to exist, the only way to realize such responsiveness is by keeping stocks. Under a responsive MPS, the safety stock level must be sufficient to realize the target service level in the presence of the variance in level of the MPS and the variance in actual period demand $d(t)$ minus the target sales $u(t, 0)$. Let $V(s)$ denote the variance in the **market-optimal Sales level** and let $V(d, u)$ denote the variance in $d(t) - u(t, 0)$. The stacked lead time of the production system is L periods. Assuming that both the target sales level $s(t)$ and $d(t) - u(t, 0)$ are not (auto)correlated, the safety stock must be able to realize the required service level in the presence of demand uncertainty with variance $L \cdot V(d, u) + L^2 \cdot V(s)$.

With a conservative MPS policy (referred to in MRP-2 parlance as a frozen MPS), CPS protects the production system and only requires it to respond to new information about the market changes in the most upstream phase; CPS does not respond to the new market information over the upcoming $L - 1$ periods and keeps also the sales targets for these future periods fixed. Note that this assumes strong control over sales level in terms of volume and timing. Such strong control will often not exist and be not cost effective if responsiveness in the production system can be achieved at low costs. Responsiveness costs may differ for different parts of the production system. If using SBS, responsiveness is created with stocks. Stock keeping costs will increase from upstream production phases to downstream production phases, in particular in divergent material flows. Thus, for a particular production situation it might be optimal to have a conservative part up to period $\xi = \rho$ in the MPS and have a responsive part from $\rho + 1$ to $L - 1$. The MPS numbers would then be $u(t, \xi) = s(t - \rho + \xi - 1)$ for $\xi = 0, \ldots, \rho$, and $u(t, \xi) = s(t)$ for $\xi = \rho + 1, \ldots, L - 1$.

The value of the delay ρ can be chosen such that it coincides with the start of a phase in the production system, thereby creating a natural position for keeping the resulting safety stock. The value of the **sales level safety stock** at that point must be sufficient to cope with a **demand level** variance equal to $(L - \rho)^2 V(s)$.

The Sales Level Decoupling Point

The MRP-2 control framework has been initially developed for situations where products are made to stock, and it still reflects this type of production strategy. Safety stocks of end-products are used to achieve a customer order service level under stochastic demand. For obvious reasons, this position in the production-sales system is referred to as the Customer Order Decoupling Point. From the analysis above, we may conclude that in production-to-stock systems that operate in dynamic stochastic markets, and operate with production processes with little flexibility, a need may exist for the creation of a **sales level decoupling point**, that is positioned upstream from the customer order decoupling point. (It is noteworthy in this context that in the MRP-2 control framework, when used for convergent divergent material flows, the MPS can be split up into Final Assembly Schedule part and a supply chain part, with the final assembly schedule as the conservative part of the MPS and the supply chain part as the responsive part. The reason for this split is that there is much less production flexibility in the final assembly process than exists the processes that produce subassemblies and component.)

6 SBS and CPS adaptions in the MRP-2 control framework

The SBS system can be readily applied in situations where market conditions are stationary and the MPS output target would therefore be flat and constant, resulting in stationary stochastic demand. The only change needed would be at the CPS level. CPS must be willing to leave the control of the stock of end-products to Material Control and specify long term service requirement that Materials Control needs to achieve. Using SBS, MC can control base stocks to achieve the service target on the condition that the MPS specified by CPS indeed represents the expected demand per period.

However, if the market is dynamic, CPS will take actions and wants to be in control of supply by changing the MPS. We should keep in mind that the actions taken by CPS may not be optimal or may even sometimes be wrong. When analyzing empirical MPS data from over periods of more than a year, and comparing these data with actual demand, it is often observed that successive MPS's are inconsistent and have almost zero predictive value at the period level. The only valuable signal could be in the **level** of the MPS. Nevertheless, since CPS is organizationally superior to MC, MC must take the MPS serious and act on it. If CPS follows a responsive MPS policy, Materials Control can each period take the new MPS data as new demand level which can be input to the SBS system. Thus, to **function** in an MRP-2 control framework under a responsive MPS policy, the SBS system should be adapted to be able to operate with a dynamic demand level.

SBS coordinates the materials flow to achieve specific service level for end-products. Thus, for SBS to be used in the MRP-2 control framework, CPS should refrain from specifying end-product stock norms and instead specify end-product service levels. SBS then can control the base stocks, taking into account the variance in the demand relative to $u(t,0)$ and the variance in the demand level, given by $V(d,u)$ and $V(s)$. With a stacked production lead time of L periods, under a responsive MPS policy, Materials Control using SBS for materials coordination is faced with demand uncertainty given by $L \cdot V(d,u) + L^2 \cdot V(s)$.

Under a completely conservative MPS policy, the SBS system needs to be adapted to operate on a varying periodic demand, the average value of which is exactly predicted in the MPS. The safety stock needed under a conservative MPS policy must be sufficient to cope with demand uncertainty with variance $L \cdot V(d,u)$, on the condition that $E(d(t)) = E(u(t,0))$, that is, sales will be able to sell on average the periodic sales targets in the MPS.

If $0 < \rho < L-1$, thus CPS uses a partially frozen MPS policy, the situation is much more complex since a new decoupling point must be introduced. We expect that this would require a redesign of SBS.

In summary, we expect that the SBS materials coordination system can be technically applied in the MRP-2 control framework under a completely conservative and under a completely responsive MPS policy. However, the use of SBS for coordinating the materials flow also requires organizational alignments. Materials Control must accept the MPS as leading, even if analysis of MPS demand prediction performance would show no or low prediction performance. Such results would suggest that using a stationary demand forecast is to be preferred. However, this would disconnect Materials Control from Coordination of Production and Sales, and would not be acceptable for CPS, since CPS is leading MC. Thus, organizational alignment requirements imply that Materials Control takes the MPS as a given and acts upon it. To the extent that the MPS contains predictive information with respect to demand, this improves the overall performance. However, following the changes in demand level requires higher safety stocks. How much safety stock would be needed in this control setting needs further investigation.

If Materials Control uses SBS for materials coordination, CPS needs to refrain from setting explicit safety stock norms and must specify a target service level for end-product demand. To be able to set an optimal service level, CPS must know the costs of safety stock resulting from a specific service level and the costs of being out of stock. The safety stock needed to achieve a specific service level results from the SBS base stock calculations and therefore can be provided by Materials Control. This allows CPS to calculate the optimal service level.

7 Conclusions

In this note we have investigated the question if and how the SBS system for materials coordination can be applied in a production control environment where control is based on concepts from the MRP-2 control framework, to be found in most widely used Enterprise Resource Planning systems. Since the MRP-2 control framework is an established way of thinking about how to control production and material flows, we looked for the modifications necessary to seamlessly embed the SBS system in such a framework, thereby replacing the MRP-1 system. We have discussed the overarching control function Coordination of Production and Sales and have analyzed the MPS functionality as device for coordinating production and sales. We have discerned an MPS policy parameter that characterizes the responsiveness of the MPS to new market information and derived the need for a new type of decoupling point, coined the Demand Level Decoupling Point, if the MPS policy is neither completely responsive nor completely conservative, but something in between.

We find that the SBS system can be directly applied if demand is stationary and stochastic. Only Coordination of Production and Sales needs a change and should specify a service level instead of a safety stock level. However, if product market is dynamic, the MPS will reflect these dynamics. If the dynamic MPS policy is either completely responsive or completely conservative, the SBS system can be applied but must be adapted to allow for a periodic update of the expected demand level, since the MPS level will periodically change, and materials coordination must follow expected demand. Again, Coordination of Production and Sales must abstain from specifying safety stock norms for end-products; instead, it must specify a service target and leave it to Materials Control to realize this service target. If the MPS policy is neither completely responsive nor completely conservative, we expect that the SBS system cannot be readily applied but will need a deeper revision.

We estimate that the changes required under a responsive or conservative MPS are quite easy to implement. As for the materials coordination system, an important consequence of these changes would be that the base stocks to achieve the service target will be higher than under a stationary stochastic demand assumption. This could be justified if expected demand indeed gradually changes over time and CPS is able to capture these changes in advance. If, however demand is stationary, and the MPS level is not stationary, it is obviously not optimal to follow the MPS. Nevertheless, if the people that run the CPS system are convinced of the rationality of changing the MPS level, following these changes is the price to pay for being embedded in the control system. The price to pay is an increase in the base stocks, the size of which could be quite easy to determine.

Two final notes of caution are warranted at this place. First, the SBS system does not use shop floor production flexibility which could be used to speed up or slow down the materials flow. If such production flexibility exists at reasonable costs, materials control might prefer a control system that exploits this flexibility. In such a situation, convincing management to apply SBS requires a demonstration of lower costs of flexibility through safety stocks that through production flow changes. Second, the Bill-of-materials processer underlying MRP-1 also facilitates product management in the timing of new product introductions and engineering changes in existing products. Thus, the SBS logic probably must in some way be 'superimposed' on the Bill-of-materials processer in MRP-2.

Ton's moments

Ivo Adan

Abstract: *In this short note I briefly elaborate on moments with Ton, socially and pleasant, and also on the many more moments of Ton, creative and scientifically inspiring.*

> Nous ne nous souvenons pas des jours,
> nous nous souvenons des moments
>
> *Cesare Pavese*

1 Just a moment

I remember many occasions, while I was working in my office, concentrating on some problems, that someone was softly knocking on my door, and entering with the request 'Do you have a moment?'. Often Ton was this person. Of course, I could not resist saying 'Of course, come in!', even though I knew that 'just a moment with Ton' could easily and most often extend to half an hour, an hour or even longer. These moments were always inspiring and enjoyable. Usually starting with some interesting problem, but then the conversation could drift in any direction, touching topics of politics, university politics, industry, teaching, mathematics, simulation, and a favorite topic, 'empirically valid models'. Indeed, in our field of Operations Management (like in physics), developing models that are validated by real data is extremely important, since we need these models to describe, analyze, understand and predict real phenomena. During our discussions

I also learned many things about supply chain models and the use of sample path arguments to derive elegant and generic optimality properties, i.e., Newsvendor equations appear everywhere. We shared our passion for mathematical modeling and analysis, and in particular, our fascination for the beauty of probability. On one occasion, I told Ton that I was passing a (transparent) classroom in the Atlas building, where a teacher was proving, through complicated and lengthy calculations, the following equality for the sum of the sizes of all subsets of a set $V = \{1, \ldots, n\}$,

$$\sum_{S \subseteq V} |S| = 2^{n-1} \cdot n, \tag{7}$$

where $|S|$ denotes the number of elements in subset S. We recognized that an elegant and simple proof of (7) can be provided by probability. Each subset S can be characterized by a vector $x = (x_1, \ldots, x_n)$ where $x_i = 1$ if $i \in S$ and 0 otherwise. Let $X_1, \ldots, X_n$ be a sequence of i.i.d. random variables where $P(X_i = 1) = 1 - P(X_i = 0) = \frac{1}{2}$. Clearly, each realization of the random vector $X = (X_1, \ldots, X_n)$ characterizes a subset of V, and the set of all realizations yields the set of all subsets of V. The mean number of elements in such a random subset is

$$E(X_1 + \cdots + X_n) = E(X_1) + \cdots + E(X_n) = \frac{1}{2} + \cdots + \frac{1}{2} = \frac{1}{2} \cdot n.$$

Noting that

$$\begin{aligned} E(X_1 + \cdots + X_n) &= \sum_{x \in \{0,1\}^n} P(X = x)(x_1 + \cdots + x_n) \\ &= \sum_{x \in \{0,1\}^n} \left(\frac{1}{2}\right)^n (x_1 + \cdots + x_n) = \left(\frac{1}{2}\right)^n \sum_{S \subseteq V} |S| \end{aligned}$$

proves equality (7). Another elegant argument was provided by Laura Sprenkels, currently a PhD student in OPAC. She observed that, by symmetry,

$$\sum_{S \subseteq V} |S| = \sum_{S \subseteq V} |V \setminus S|.$$

Hence,

$$2 \sum_{S \subseteq V} |S| = \sum_{S \subseteq V} |S| + \sum_{S \subseteq V} |V \setminus S| = \sum_{S \subseteq V} (|S| + |V \setminus S|) = \sum_{S \subseteq V} n = 2^n \cdot n,$$

which also yields (7). In the next section I proceed to discuss infinitely many more moments of Ton.

2 Many more moments

In this section, we consider the standard $G/G/1$ queue. Customers arrive according to a renewal process with inter-arrival time distribution $F_A(\cdot)$. The service-time distribution is given by $F_B(\cdot)$. The service times of customers are assumed to be mutually independent and independent of the arrival process. Customers are served in order of arrival by a single server. We assume that the queue is stable, that is, $\rho := E(B)/E(A) < 1$, where the random variables A and B are distributed according to $F_A(\cdot)$ and $F_B(\cdot)$, respectively.

For $n = 0, 1, 2, \ldots$, let the random variables A_n denote the inter-arrival time between the $(n-1)$th and nth customer, the random variables B_n and W_n denote the service time and the waiting time of the nth customer, respectively. Then it is readily seen that the following equation holds for the waiting times

$$W_n = \max\{0, W_{n-1} + B_{n-1} - A_n\}, \quad n = 1, 2, \ldots \tag{8}$$

This equation is commonly referred to as Lindley's equation. As n tends to infinity, the distribution of W_n converges to a limiting distribution. Despite the simplicity of Lindley's equation, an exact solution for the limiting distribution of the waiting time, or its moments, is only possible in special cases of $F_A(\cdot)$ and $F_B(\cdot)$. Hence, this justifies the use of approximations, and Ton de Kok (1989) had the bright idea to use Lindley's equation to iteratively approximate the limiting probability of waiting and the first two moments of the waiting time. By Lindley's equation (8), the waiting time distribution of the nth customer can be expressed in terms of the waiting time distribution of the $(n-1)$th customer. From (8) we obtain the following expressions for $P(W_n > 0)$, $E(W_n)$ and $E(W_n^2)$,

$$\begin{aligned}
P(W_n > 0) &= \int_0^\infty (1 - F_{W_{n-1}+B}(t))dF_A(t), \\
E(W_n) &= \int_0^\infty \int_t^\infty (y-t)dF_{W_{n-1}+B}(y)dF_A(t), \\
E(W_n^2) &= \int_0^\infty \int_t^\infty (y-t)^2 dF_{W_{n-1}+B}(y)dF_A(t),
\end{aligned}$$

where $F_{W_{n-1}+B}(\cdot)$ denotes the distribution of $W_{n-1} + B$. This distribution is unknown, but if estimates for the first two-moments of W_{n-1} would be available, then we can calculate the first two moments of $W_{n-1} + B$ (since these random variables are independent) and fit *tractable* (continuous) distributions $\tilde{F}_{W_{n-1}+B}(\cdot)$ and $\tilde{F_A}(\cdot)$ to the distributions of $W_{n-1} + B$ and A respectively, by matching the first two moments. A common recipe for fitting a tractable continuous distribution on the mean, $E(X)$, and the coefficient of variation, c_X, of a given non-negative random variable X is the following (see e.g. Tijms 1986). In case $0 < c_X < 1$, one fits a mixture of Erlang distributions with the same scale parameter. Otherwise, in case $c_X \geq 1$, one fits a hyper-exponential distribution with balanced means. These fitted distributions are tractable in the sense that the above integrals for $E(W_n)$ and $E(W_n^2)$ can be easily calculated. Hence, we obtain estimates for the first two moments of W_n, and we can repeat this procedure to obtain estimates for W_{n+1}, W_{n+2} and so on, until we reach stationarity (i.e., convergence). This iteration scheme can be initiated by any choice for $E(W_0)$ and $E(W_0^2)$, e.g., by assuming that the 0th customer arrives in an empty system, so $W_0 = 0$. The resulting scheme can be summarized as follows (see de Kok 1989).

Moment-iteration algorithm

Step 0 (Initialization): Set $E(W_0) = E(W_0^2) = 0$.
Step 1 (Iteration): Compute the first two moments of W_{n-1} and B,

$$\begin{aligned} E(W_{n-1} + B) &= E(W_{n-1}) + E(B), \\ E((W_{n-1} + B)^2) &= E(W_{n-1}^2) + 2E(W_{n-1})E(B) + E(B^2). \end{aligned}$$

Fit tractable distributions $\tilde{F}_{W_{n-1}+B}(\cdot)$ and $\tilde{F}_A(\cdot)$ to the distributions of $W_{n-1}+B$ and A respectively, by matching the first two moments, and compute

$$\begin{aligned} E(W_n) &= \int_0^\infty \int_t^\infty (y-t) d\tilde{F}_{W_{n-1}+B}(y) d\tilde{F}_A(t), \\ E(W_n^2) &= \int_0^\infty \int_t^\infty (y-t)^2 d\tilde{F}_{W_{n-1}+B}(y) d\tilde{F}_A(t). \end{aligned}$$

Step 2 (Convergence): If $|E(W_n) - E(W_{n-1})| \leq \epsilon$ and $|E(W_n^2) - E(W_{n-1}^2)| \leq \epsilon$ (for some $\epsilon > 0$), then stop, and otherwise repeat Step 1.
Step 3 (Stop): Approximate the first two moments of the limiting distribution of the waiting time by $E(W_n)$ and $E(W_n^2)$, and probability of waiting by

$$P(W_n > 0) = \int_0^\infty (1 - \tilde{F}_{W_{n-1}+B}(t)) d\tilde{F}_A(t)$$

This algorithm is simple, elegant, easy to implement, and effective: the approximations are good quality and robust. This work inspired the analysis in Adan et al. (1995). In that paper, we consider the $D/G/1$ queue with a *discrete* service-time distribution, which arises, for example, in the analysis of the fixed-cycle traffic light Darroch (1964) and in the periodic review (R, S) inventory system with finite production capacity de Kok (1989). To compute the waiting-time characteristics of the discrete-time $D/G/1$ queue, we applied Ton's moment-iteration algorithm. However, to fit a tractable distribution to the first two moments of the discrete distribution of $W_{n-1} + B$, it is more natural (and more accurate) to also fit a discrete distribution, instead of a continuous one. Hence, this spurred the development of a discrete analogue of the simple fitting procedure for continuous contributions (as described above). It turned out that in order to fit all feasible values of $E(X)$ and c_X in the discrete case, four classes of distributions are required, instead of only two, namely mixed-binomial, Poison, mixed negative-binomial and mixed geometric distributions. Use of this recipe to fit a discrete distribution to the first two moments in Ton's moment iteration algorithm produced accurate approximations for the waiting time characteristics of the discrete-time $D/G/1$ queue. This fitting recipe also found its way in many other applications where discrete distributions are relevant. It is interesting to note that Lindley's equation is subtle in the sense that only a change of sign of W_{n-1} in (8), so

$$W_n = \max\{0, B_{n-1} - A_n - W_{n-1}\}, \quad n = 1, 2, \ldots,$$

turns the equation into a tractable one. The above equation (with the minus sign for W_{n-1}) naturally arises from the performance analysis of a system consisting of two parallel carousels operated by a single picker (Vlasiou et al. 2004).

3 Series of moments

This section presents another application of Ton's moment iteration idea. We consider a serial queueing system without waiting space between successive servers, which was motivated by the modeling and analysis of an assembly line in the automotive industry (Frenken 2007). This example illustrates that research quite often stands on the shoulders of giants, and a giant is Ton! The assembly line is the SIBO area (Sides In Bodies Out) in the body shop of a car manufacturer. The assembly line is schematically shown in Fig. 2. It consists of 35 workstations, where the products are transported on carriers. Each workstation can hold one product carrier and represents an operator-handled machine, a robot-handled machine, or a buffer place. The processing times collected from each of the machines are *effective* processing times (Hopp and Spearman 2008), that is, raw processing times including all realities on the store floor (breakdowns, setups, etc.). The processing times of the buffer places are the transport times to the next position

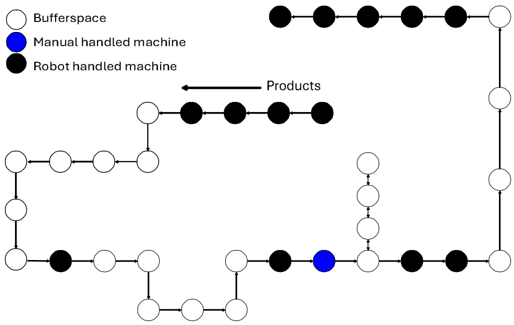

Figure 2: SIBO area consisting of 35 workstations.

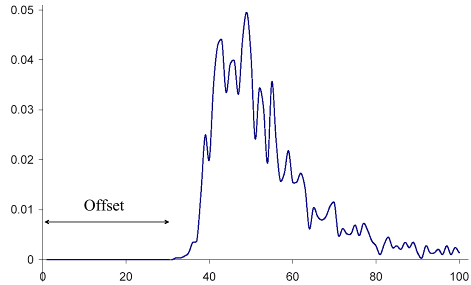

Figure 3: Empirical probability distribution of effective processing time with fixed offset

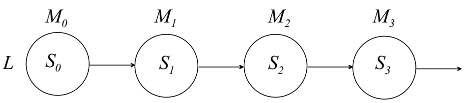

Figure 4: Serial line L with machines M_0, M_1, M_2, M_3.

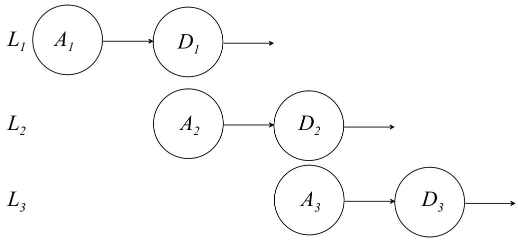

Figure 5: Decomposition of L into two-machine subsystems L_1, L_2, L_3.

in the line. A key feature is that processing times have a fixed offset (depending on the workstation), see Fig. 3. The purpose of the analysis was to develop an efficient method to estimate the throughput capacity of the assembly line.

The assembly line is modeled as a serial line L with N machines $M_0, \ldots, M_{N-1}$ and no buffer space between the machines, see Fig. 4. Hence, the machines are tightly coupled. This implies that, when a machine M_i finishes processing a product and the downstream machine M_{i+1} is still processing, machine i will be blocked and has to wait until machine M_{i+1} becomes available before it can transfer the product. So, the blocking protocol is *block-after-service*. The last machine M_{N-1} is never blocked and, since we are interested in the throughput capacity of the line, we assume that machine M_0 is never starved (that is, there are always products to be processed). The random variable S_i denotes the effective processing time of machine M_i, which has offset o_i. So $S_i - o_i$ is the random part of the processing time at machine M_i. Note that the buffer places in the assembly line in Fig. 2 are modeled as machines, the processing times of which are the transportation times.

To develop an approximation for the throughput, the serial line is decomposed into subsystems L_i consisting of machines M_{i-1} and M_i (cf. Gershwin 1987). The decomposition of L is shown in Fig. 5. The subsystems, of course, do not operate independently. This dependence is reflected in the processing times A_i and D_i of M_{i-1} and M_i in subsystem L_i. To define these processing times, we note that machine M_{i-1} may have been starved before it starts processing a product and machine M_i can become blocked after finishing a product. This motivates the following definition of processing times in subsystem L_i. Let $A_{i,n}$ and $D_{i,n}$ denote the nth processing time of machine M_{i-1} and M_i, respectively. Then:

- $A_{i,n}$ is the nth effective processing time $S_{i-1,n}$ of M_{i-1} plus the starvation time before the start of this processing time,
- $D_{i,n}$ is the nth effective processing time $S_{i,n}$ of M_i plus the blocking time of M_i after the end of this processing time.

If we define a cycle as the time that elapses between two successive product transfers from M_{i-1} to M_i, then the nth cycle of L_i can be expressed as

$$C_{i,n} = \max\{A_{i,n}, D_{i,n-1}\}. \tag{9}$$

To derive an expression for $A_{i,n}$, we note that $\max\{A_{i-1,n} - D_{i-1,n-1}, 0\}$ is the starvation time of M_{i-1} before the start of $S_{i-1,n}$, and thus

$$A_{i,n} = \max\{A_{i-1,n} - D_{i-1,n-1}, 0\} + S_{i-1,n}. \tag{10}$$

The blocking time of M_i after the end of $S_{i,n}$ is $\max\{D_{i+1,n-1} - A_{i+1,n}, 0\}$. Hence, we can write

$$D_{i,n} = S_{i,n} + \max\{D_{i+1,n-1} - A_{i+1,n}, 0\}.$$

This relation is symmetric to (10). Note that $A_{i+1,n}$ depends on $S_{i,n}$, since by (10), we have $A_{i+1,n} = \max\{A_{i,n} - D_{i,n-1}, 0\} + S_{i,n}$. Substitution of this equation in the above relation for $D_{i,n}$ yields

$$\begin{aligned} D_{i,n} &= S_{i,n} + \max\{D_{i+1,n-1} - \max\{A_{i,n} - D_{i,n-1}, 0\} - S_{i,n}, 0\} \\ &= \max\{D_{i+1,n-1} - \max\{A_{i,n} - D_{i,n-1}, 0\}, S_{i,n}\}. \end{aligned} \tag{11}$$

We assume that initially, the production line is empty, implying that $D_{i,0} = 0$ for all i. Since M_0 is never starved and M_{N-1} is never blocked, we have $A_{0,n} = S_{0,n}$ and $D_{N-1,n} = S_{N-1,n}$ for all n.

The random variables A_i and D_i in Fig. 5 refer to the steady-state versions of $A_{i,n}$ and $D_{i,n}$ as n tends to infinity. We aim to estimate the steady-state throughput

$$TH = \lim_{n\to\infty} \frac{1}{E(C_{N-1,n})}.$$

Now we are going to apply Ton's idea to use the sample path relations (10) and (11) to iteratively approximate the throughput TH. This results in the following algorithm.

Step 0 (Initialization): $D_{i,0} = 0$ for all $i = 0, 1, \ldots, N-1$.
Step 1 (Iteration): For $i = 1, 2, \ldots, N-1$,

(i) Fit tractable distributions to the distributions of the random variables $A_{i-1,n}$, $D_{i-1,n-1}$, $D_{i,n-1}$ and $D_{i+1,n-1}$ by matching the first two moments of their *random parts*.

(ii) Compute the first two moments of $A_{i,n}$ and $D_{i,n}$ from (10) and (11) by acting as if $A_{i-1,n}$, $D_{i-1,n-1}$, $D_{i,n-1}$ and $D_{i+1,n-1}$ are *independent*.

(iii) Compute $E(C_{i,n})$ from (9).

Step 2 (Convergence): If $|E(C_{i,n}) - E(C_{i,n-1})| \leq \epsilon$ for all i (for some $\epsilon > 0$), then stop, and otherwise repeat Step 1.
Step 3 (Stop): Approximate the steady-state throughput TH by $1/E(C_{N-1,n})$.

This algorithm has been evaluated for a large (synthetic) test set of more than 2500 cases, including production lines with up to 40 workstations, squared coefficients of variation ranging from 0.2 to 2, offsets from 0% to 90% of the mean effective processing time, balanced and unbalanced lines (for details, see Frenken 2007). The overall average error of the steady-state throughput estimate appeared to be less than 4.5%, and less than 1.5% for cases without offset in effective processing times. But the proof of the pudding is in the eating, so we applied this algorithm to the SIBO area. The empirical distribution of the effective process time in the operator-handled machine in the SIBO area (i.e., workstation 20 in Fig. 2) and in one of the robot-handled machines is shown in Figs. 6 and 7. The distribution of the manual station is spread out with no clear peak, the one of the robotic station has a distinctive peak with occasionally some higher values.

The offset, mean and squared coefficient of variation of effective processing times in the SIBO area are summarized in Table 2. The variability of the effective

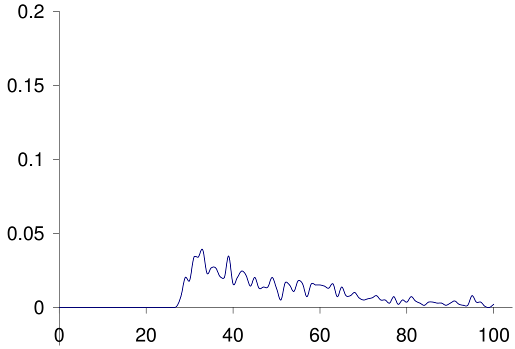

Figure 6: Empirical probability distribution of processing time of operator handled machine.

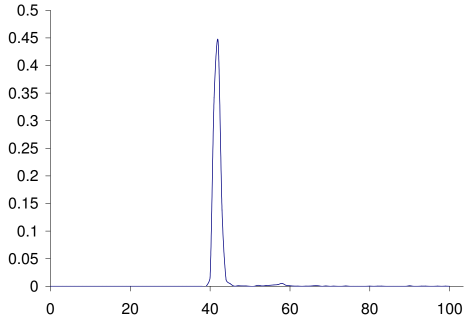

Figure 7: Empirical probability distribution of processing time of robot handled machine.

machine	1	2	3	4	5-11	12	13-18	19	20	21-23	24
offset	66	63	53	33	8.3	58	8	51	29	0	7
mean	70	71	59	49	8.3	67	8	59	64	0	7
cv^2	.84	2.8	.48	.15	0	.13	0	4.5	.41	0	0
machine	25	26	27	28	29	30	31	32	33	34	35
offset	69	67	20	7	8	22	42	61	55	60	21
mean	73	69	20	7	8	22	51	67	57	65	37
c_Z^2	.13	.04	0	0	0	0	8.0	.67	.11	.03	6.3

Table 2: Offset, mean and squared coefficient of variation of the effective processing times of the workstations in the SIBO area.

processing times looks reasonable, but one has to realize that the variability of the random part of the effective processing times is usually much higher. For example, the squared coefficient of variation of the effective processing time in workstation 1 is .84, but the squared coefficient of variation of the random part of this effective processing time is $.84 \times (70/(70-66))^2 = 257$!

For the SIBO area, we constructed four cases. The first case consists of the last 5 workstations, so stations 31 up to 35. The second case consists of workstations 25 up to 35, the third case consists of workstations 19 up to 26 (so this case includes the operator-handled machine), and the fourth case consists of all 35 workstations. The results are shown in Table 3. The performance of the algorithm is good across all four cases, so this is a first promising step in the empirical validation of the proposed model and algorithm. I think Ton will be satisfied!

Case	1	2	3	4
Error % in Throughput	+1.57	+4.77	+1.12	+1.13

Table 3: Error % in estimate for steady-state throughput for cases 1-4.

Ton's Work on Inventory Systems with Discrete Product Modeling

Geert-Jan van Houtum

When modeling an inventory system, we have to choose whether the considered product is modeled as a discrete or a continuous product. Most of the research of Ton de Kok has been on continuous-product inventory systems, but he also worked on discrete-product inventory systems. In this chapter, I write about that part of his research.

1 Introduction

When modeling an inventory system, many modeling choices have to be made. One of the choices concerns the modeling of the product, which can be assumed to be discrete or continuous. Similarly, time can be modeled as continuous or discrete. In the standard EOQ model, both the time and the product are continuous. In the serial multi-echelon inventory system of Clark and Scarf (1960), time is discrete (i.e., time consists of periods of length one) and the product is continuous. In the well-known METRIC model of Sherbrooke (1968) for a multi-echelon spare parts inventory system, the time is continuous and the product is discrete.

When you have low-demand products, such as spare parts, library books, and airplanes, it seems natural to choose discrete-product modeling. Continuous-product modeling would not make sense in that case. Generally, we see that researchers in the field of inventory management prefer to model product(s) as continuous. This often smoothens the analysis. For example, the analysis of the

EOQ model and the resulting EOQ formula would be slightly less clean when a discrete product would be assumed. In the multi-echelon model of Clark and Scarf (1960), due to the continuous product assumption, a convex one-period cost function is obtained on a continuous domain, and this convexity propagates to other relevant cost functions with continuous domains. This is a key point in the proof of the optimality of basestock policies. Continuous demand modeling is generally also more practical for calculations (see also Axsäter (2006), Chapter 5).

Regarding time, it seems that researchers often choose for discrete time rather than continuous time. This choice is natural when replenishment orders are always placed at the beginning of a week or month, or when demand data have been collected per week or month. In addition, the choice for a discrete time facilitates the coordination of orders for different products. It is also convenient when you want to apply stochastic dynamic programming in order to study optimal policy structures. Nevertheless, we see also streams of literature with continuous-time modeling like all kinds of EOQ models and spare parts inventory models. In multi-echelon spare parts systems with a divergent structure (like METRIC), the combination of basestock control, Poisson demand processes at the most downstream inventory locations, and continuous time (and review) leads to the property that demand processes at all locations are Poisson processes. This facilitates the exact analysis of a given basestock policy (see e.g. Rustenburg et al. (2003)). For more reflections on discrete- versus continuous-time modeling, see Section 3.2.2 of Axsäter (2006).

A lot of the research of Ton de Kok is in the area of multi-echelon inventory systems and there continuous products and discrete time are assumed. There are two exceptions: (i) The one-warehouse, multi-retailer inventory systems studied in the PhD project of Mustafa Doğru; (ii) Research on spare parts inventory systems. The rest of this chapter is devoted to these two exceptions. In both exceptions, products are assumed to be discrete, while the time is discrete in exception (i) and continuous in exception (ii).

2 One-warehouse, multi-retailer systems

Ton's first inventory paper appeared in 1984 (see De Kok et al. (1984)). I worked on my first inventory research project for my master thesis as part of my applied mathematics program at TU/e in 1989-1990. My supervisor was Henk Zijm, who was a part-time professor at TU/e. His main job was at CQM, where he was a direct colleague of Ton. Henk and Ton interacted a lot on inventory problems, or more generally OR problems. Also other colleagues were involved in these discussions (see the chapter of Mynt Zijlstra in this liber amicorum). If I remember it correctly, then Ton and I never met each other in that time. We met each other probably for the first time in 1991 when Ton was appointed as part-time professor at the Dept. Mathematics and Computer Science, where I then worked as a PhD student on a queueing project (with Wessels, Zijm and Adan as supervisors).

In the years 1995-1996, we had a kind of first collaboration. On invitation of the *European Journal of Operational Research*, we worked on two complementary review papers on multi-echelon production-inventory systems:

- A paper by van Houtum, Inderfurth, and Zijm on "Materials coordination in stochastic multi-echelon systems". This paper was focused on pure cost models, where the costs consist of inventory holding costs at different stages and backordering costs (see Van Houtum et al. (1996)).
- A paper by Diks, de Kok and Lagodimos on "Multi-echelon systems: a service measure perspective". In this paper, the focus was on models with service level constraints instead of backordering costs (see Diks et al. (1996)).

It was not a coincidence that Ton was involved in the second paper. He always had a strong preference for a direct modeling of service constraints (see also the chapter of Henk Zijm in this liber amicorum).

In 2001, Ton and I started a real collaboration when we jointly supervised Mustafa Doğru, who graduated in 2005 with a PhD thesis on "Optimal Control of One-Warehouse Multi-Retailer Systems: An Assessment of the Balance Assumption". When analyzing multi-echelon inventory systems with a distribution structure, often the so-called *balance assumption* is made in order to obtain a decomposition results as for the serial Clark-Scarf system (see Clark and Scarf (1960)). However, how bad it was to make this balance assumption was never well studied at that time. Mustafa studied the effect of the balance assumption by comparing: (i) the optimal policy for the true system without the balance assumption; (ii) the heuristic policy obtained after making the balance assumption. Mustafa switched from the common continuous-product assumption to the discrete-product assumption so that the optimal policy could be computed relatively easily for small and medium-size systems. The relative gap between the costs of the heuristic and optimal policy in the true system indicates whether the heuristic policy performs well and hence how strong the effect is of making the balance assumption.

In a test bed with symmetric problem instances (i.e., identical parameters for all local warehouses), the relative gap was in most instances small, and in almost all

Figure 8: Ton as an active participant at one the sessions of the 20-th International Symposium on Inventories in Budapest, 2018.

instances below 10%. However, comparing both solutions in a large test bed with asymmetric problem instances, showed that the relative gap was more than 10% in a significant fraction of instances, and three instances were identified with a gap of more than 100%. In the latter three instances, there were two local warehouses: one local warehouse with a high mean demand per period, a low coefficient of variation, and a high penalty cost parameter and another local warehouse with a low mean demand per period, a high coefficient of variation, and a low penalty cost parameter. This was a clear commonality between the three problem instances with a gap of more than 100%. In general, there was not a clear picture of which parameter combinations lead to the larger relative gaps. For more details on this research, see Chapter 4 of Doğru (2005).

3 Spare parts inventory systems

Ton de Kok is well-known for his research on supply chains for new products, say. However, he also had a significant contribution in the area of spare parts supply chains. In the 1990-s, he was a board member of the Service Logistics Forum, which in that time had the name Parts Business Forum. In that period, he supervised Jos Verrijdt, who was Ton's first PhD student and worked on spare parts inventory control. In total, Ton supervised more than 40 PhD theses as first or second supervisor. Seven of them were on spare parts:

- Jos Verrijdt - Design and Control of Service Part Distribution Systems (1997; supervisors: de Kok and Theeuwes)
- Jan Willem Rustenburg - A System Approach to Budget-Constrained Spare Parts Management (2000; supervisors: Zijm, de Kok and van Houtum)
- Bram Kranenburg - Spare Parts Inventory Control under System Availability Constraints (2006; supervisors: de Kok and van Houtum)
- Ingrid Vliegen - Integrated Planning for Service Tools and Spare Parts for Capital Goods (2009; supervisors: van Houtum and de Kok)
- Kurtuluş Öner - Optimal Reliability and Upgrading Decisions for Capital Goods (2010; supervisors: van Houtum and de Kok)
- Joachim Arts - Spare Parts Planning and Control for Maintenance Operations (2013; supervisors: van Houtum and de Kok)
- Zhao Kang - Robust Spare Parts Inventory Management (2025; supervisors: Basten, de Kok, and Marandi)

Jos Verrijdt was the second PhD student of Ton (based on graduation date) and the first one for whom Ton acted as first promotor. In the preface, Jos writes: "Ton de Kok, my advisor during the last five years in which I worked on my master thesis and PhD thesis, provided invaluable support. He generated *a never-ending stream of ideas, suggestions, and comments* that contributed much to the contents of this thesis." This point of a never-ending stream of ideas, suggestions, and comments is probably recognized by many PhD students. It is a strength of Ton and it was intertwined with his enthusiasm for many research projects. His mathematical and queueing background also helped a lot to see the possibilities for extensions and connections with other inventory problems. Ton always had a rich international network of whom many people have benefited. Jos Verrijdt had the opportunity to collaborate for one of his subprojects with Patrik Alftredsson of the Royal Institute of Technology in Stockholm. Together they wrote a very nice paper that was published in *Management Science* (see Alfredsson and Verrijdt (1999))

Jan Willem Rustenburg was a navy officer who executed his PhD project at the Royal Netherlands Navy Dockyard in Den Helder. Ton got involved in the last

Figure 9: Ton is known for generating streams of ideas and suggestions. At this picture, he consumes a stream of information from a guide during a trip to Volendam with the OPAC group.

phase of this project as second promotor. The projects of Bram Kranenburg and Ingrid Vliegen were executed in close collaboration with ASML. During many years, Ton and I collaborated with ASML, where Ton focused on the supply chain for new machines and I on the after sales supply chain. The project of Kranenburg resulted in the development of a new spare parts planning concept that was implemented in 2005. After further developments, it resulted in the current SPARTAN system at ASML.

Kurtuluş Öner studied the effect of reliability and redundancy decisions on the total spare parts costs for capital goods. That work was strongly motivated by the practice at Philips. He also looked at so-called upgrading decisions for redesigned components, which was motivated by the practice at ASML. Joachim Arts wrote a thesis on the maintenance and spare parts management for trains. He collaborated closely with Nedtrain, the maintenance organization of the Dutch

Railways. The last spare parts thesis by Zhao Kang has been strongly inspired by a problem of ASML in the initial phase after the introduction of a new machine in the field. In that phase, you have limited information on the failure rates of the components, while you do have to take inventory decisions for spare parts. A company can then decide to make rough estimates for the demand rates and apply the traditional spare parts inventory models. Zhao applied the concept of robust optimization and showed that this leads to better decisions.

Ton, it was always an honor and an immense pleasure to collaborate with you, on the topics described above, but also on many other things such as events for the European Supply Chain Forum, INFORMS events, strategy workshops of the OPAC group, the Beta Research School and of course on making fun. Thanks a lot for everything!

Analytical Planning Stability Research: No Reason for Nervousness!

Karl Inderfurth

1 Introduction

Planning stability or its counterpart, planning nervousness, is a challenging issue for many companies as well as for scientists who try to analyze and explain this phenomenon. It was, however, not just this topic that stood at the beginning of the collaboration between Ton de Kok and me. We both spent some time of practical work with a company before we started our university careers. In practice, we both were confronted with challenging problems from the field of supply chain planning which we made to a central subject of our scientific work after switching to academia. This common interest in supply chain optimization brought us together, and it is this area where one of our two joint papers (Inderfurth et al. 2001) refers to.

The other publication (De Kok and Inderfurth 1997) belongs to a different field and considers the issue that this paper is dedicated to, namely the analysis of nervousness in planning systems. The problem is old, but research in this area for a long time was only descriptive or simulation-based. This changed in the middle 1990s when the first analytical investigation was published in Inderfurth (1994). In this context, I was happy that I could draw Ton's attention to this research topic so that a collaboration started in which I could profit a lot from Ton's enormous analytic expertise. Starting with this joint research, the present paper presents an overview of the outcome of the analytical research contributions on the topic of planning stability. It describes the advantages of gaining analytical insights, but also the technical difficulties of this research approach which might have deterred many researchers from trying to analyze planning stability in this way. A matter of nervousness in tackling hard problems?

2 Planning Stability Research

Research in the area of planning stability mainly concentrates on the field of production and materials planning. Standard planning systems like MRP are generally operating in a rolling horizon framework where after each period a re-planning takes place according to updated information. Updates, in general, stem from a prolongation of the planning horizon and from forecast errors and new demand estimates. Thus, changes in the quantity or timing of planned orders are generated which create system nervousness. Planning stability research focusses on how MRP rules and policies affect the size of planning instability. In practice, different types of Period Order Quantity (POQ) or Fixed Order Quantity (FOQ) lotsizing are used in an MRP context. Under demand uncertainty these lotsizing policies are equivalent to employing stochastic inventory control rules of the (s, nQ), (s, S), or (R, S) type (see Lagodimos and Anderson 1993). Therefore, much research concerning planning nervousness refers to the impact of these rules and their parameters.

Usually, in the field of production and inventory planning the economic effect of different control rules is measured in terms of operating cost. In many cases, however, a pure cost-oriented valuation of effects is not sufficient so that an additional measure is needed. This is well-known for the objective of reaching a high service degree. If cost of material shortages can hardly be estimated, a service level is used as an additional criterion to assess the performance of a planning system. Next to service deficiencies, frequent replanning activities can harm the efficiency of planning procedures. In general, the impact of planning nervousness is much harder to value in terms of costs than the service impact. So, the level of planning stability can be viewed as the third attribute that, next to cost and service level, is a relevant performance criterion of a planning system. Accordingly, almost all scientific contributions in the field address the influence of ordering rules on the stability level isolated from cost considerations.

The majority of research papers which investigate the impact of control rules on system nervousness is based on simulation methods. This starts with a first study by Blackburn et al. (1986) and continues to recent papers like those by Atadeniz and Sridharan (2020) and Sáez et al. (2023). Many of these papers consider the impact of different standard planning procedures, others are dedicated to the analysis of specific planning tools which are designed for dampening planning instability. Simulation-based studies can cope with complex planning situations met in multi-product, multi-stage production systems or with different kinds of stochastic inputs and non-stationarities. These research contributions also allow for an analysis of the impact of various dampening methods on planning nervousness. They, however, are bound to a specific planning environment and do not allow for deeper and general insights into the interrelationship between specific control rules and the degree of planning instability generated by them. Additionally, these approaches typically use problem-specific measures for the size of planning nervousness which cannot be generalized. Thus, a comparison of the nervousness level under different planning scenarios is not possible. To overcome the shortcomings of these simulation-based procedures analytical approaches have been developed. For sake of mathematical tractability, they mainly are restricted to planning problems with a single-item, single-stage environment and to the application of basic inventory control rules. Despite these limitations they provide important insights into the creation of nervousness. These analytical approaches and their general findings will be described in the sequel.

3 Analytical Approaches

In all analytical approaches under consideration the main inventory control rules, namely the (s, nQ), (s, S), and (R, S) rule, are investigated under stochastic demand conditions. Completely different from standard stochastic inventory research which aims to find optimal control parameters for a single planning instance, nervousness research analyzes steady state conditions of two consecutive planning runs and compares ordering decisions for the same planning periods in these runs. By this way it is possible to develop closed-form formulas which describe how the level of planning nervousness in form of order deviations depends on the choice of the specific control rule and its parameters. In order to investigate the interdependence of consecutive planning instances analytically, usually some restrictions concerning the planning situation must be accepted. Only single-stage stationary inventory systems with stochastic demand and an infinite planning horizon are considered. The demand per period is assumed to be a stationary and independent random variable. The demand's expected value is used as (constant) demand forecast (denoted by D) for each planning period of a cycle. Thus, order changes for the same period from cycle to cycle are generated by deviations of the actual demand in the first period from its forecast the period before.

In the literature, we only find a very limited number of contributions which are dedicated to the described type of analysis. Chronologically, these are Inderfurth (1994), De Kok and Inderfurth (1997), Heisig (1998), Heisig (1998), Heisig and Fleischmann (2001), and Heisig (2002). These approaches differ with respect to the types of order deviations and number of compared planning periods they consider. Depending on the practical relevance of order changes for planning stability, it can be the pure deviation in order quantity from one planning run to the next that matters or, alternatively, the fact that irrespective of the quantity the setup characteristic of orders is changing. Accordingly, we differentiate the types of *quantity-oriented* and *setup-oriented* planning stability. In a temporal respect, only order deviations concerning the first period of a planning horizon might be relevant in practice so that it is *short-term* stability of orders that matters. If multiple periods play a role in the valuation and measurement of order deviations we refer to *long-term* stability.

Thus, in combination we end up with four types of stability cases which reflect the major aspects of nervousness in practical problems. For the sake of comparability of nervousness levels under different problem data and under different control rules, additionally a normalization of the nervousness metric is needed. This aspect is first discussed in Jensen (1993) where a normalized measure is defined in which the size of order deviations under concern is divided by the respective maximum deviation that can occur. In that way, all of the four stability measures introduced above are normalized such that they can only take on numerical values between 0 (minimum stability) and 1 (maximum stability). This is very similar to the normalization of service measures in inventory control where respective service levels (of different types) are limited to an interval between 0 and 1. Concerning these planning stability measures, the six literature contributions with analytical stability research present deep insights into the planning stability characteristics of the three control rules under consideration.

4 Findings

4.1 Overall Findings

The first finding refers to the impact of the specific control rule parameter which is mainly responsible for the service level that a rule provides. This is the reorder point s for the (s, nQ) and the (s, S) rule, and the reorder level S for the (R, S) rule. It turns out that these rule parameters have no impact on planning nervousness at all. This holds for both *short-term* and *long-term* planning stability instances. This property can easily be evaluated. It is due to the fact that the respective control parameters in steady state do not influence the sequence and size of order decisions. They only determine the basic inventory level of the production system. Thus, as a general finding we can secure that optimizing customer service and planning stability are no conflicting goals in this planning context.

With respect to the impact of the other control rule parameters, we must consider the specific rules and differ between *short-term* and *long-term* stability.

4.2 Short-term Stability

When applying an (s, nQ) policy, planning nervousness is only affected by the size of the standard lotsize Q. In De Kok and Inderfurth (1997) it is shown how the level of planning stability π depends on the choice of parameter Q. A closed-form solution for the so-called stability function $\pi(Q)$ is derived which can be exploited both analytically and numerically.

Starting with the results for *setup-oriented* stability, it turns out that this measure approaches 100% when the lotsize tends to be very small ($Q \to 0$) or very large ($Q \to \infty$). This result is intuitive. For $Q \to 0$ the (s, nQ) policy turns into a simple order-up-to-s policy which results in a planned and executed setup in each period of any planning cycle. On the other hand, also for $Q \to \infty$ more and more planned and executed orders within a cycle become equal because of an increasing number of periods without any setup. Inside these boundaries, $\pi(Q)$ reacts such that the planning stability is quickly decreasing when Q increases until π reaches a minimum level in the region $D \leq Q \leq 2 \cdot D$. Thus, the size of the lotsize/demand ratio is responsible for the instance of maximum planning nervousness. When Q rises further on, planning stability π increases again. The detailed shape of this unimodal stability function as well as the location of its minimum and the respective stability level depend on the properties of the stochastic demand distribution. If a very general distribution function in form of a mixture of two Erlang distributions is chosen, the values of $\pi(Q)$ can be determined numerically for various parameters of the demand distribution. A respective investigation in De Kok and Inderfurth (1997) shows that – as can be expected – the demand variance has a significant impact on the course of $\pi(Q)$. It appears that for all values of Q the level of nervousness increases if demand variability rises. To get an idea of the numerical value of the stability measure, the π value will be reported for the specific lotsize Q which triggers minimum planning stability. In this case, the size of the stability level goes down to 57% if the squared coefficient of demand variation (CV) is very high ($CV = 2.0$) and reaches 87% for low variation ($CV = 0.25$).

The respective property of the stability function $\pi(Q)$ is very different if we consider *quantity-oriented* planning stability. When deviations in the complete order quantity matter, a 100% stability cannot be reached under random demand. Surprisingly, however, this type of stability does not change with a variation of lotsize Q. The stability level only depends on the properties of the demand distribution and, specifically, on the demand variability. With respect to the CV impact, its size is close to the values reported above for the minimum level for *setup-oriented* stability. The main finding from this analysis is that *quantity-oriented* stability cannot be controlled by parameter choice under an (s, nQ) policy.

When an (s, S) rule is applied, an analogous analysis of ordering-related nervousness can be performed. Here, we find closed-form expressions for a stability function $\pi(Q_M)$ which describes how stability measure π depends on minimum lotsize Q_M (with $Q_M := S - s$). So, under the (s, S) rule the lotsize parameter is defined by the spread between reorder level and reorder point. For *setup-oriented* stability the same property of 100% stability holds for the lotsize limits $Q_M \to 0$ and $Q_M \to \infty$ like under the (s, nQ) policy. Concerning the complete course of stability function $\pi(Q_M)$, a different picture emerges. Starting from the 100% level at $Q_M \to 0$, stability π sharply declines with increasing Q_M and reaches its minimum value exactly at $Q_M = D$. A further rise of Q_M results in an upward jump of stability π followed by a monotone increase. Thus, as a specific property of the stability function for an (s, S) policy we find that a discontinuity appears at a lotsize level which equals the demand forecast. At this order size, which represents a lot-for-lot policy, we face minimum planning stability with values of only 45% for demand variability $CV = 2.0$ and 61% for $CV = 0.25$.

When we turn to *quantity-oriented* stability, we find that the stability function in the case of an (s, S) policy shows a behavior that is very different from the one under an (s, nQ) rule. Only for the lotsize limits $Q_M = Q \to 0$ and $Q_M = Q \to \infty$ the stability values for both control rules coincide. Under an (s, S) policy, however, the stability level π does not remain constant for varying lotsize choices. Instead, stability always undershoots the (s, nQ) level. The respective stability function has properties similar to those in the case of *setup-oriented* stability. That means that we find a discontinuity at the lotsize level $Q_M = D$ which also characterizes the location of minimum planning stability.

Different from the situation under reorder-point control rules, the analysis of planning stability under an (R, S) policy turns out to be quite simple. Because the replenishment cycle has a fixed length of R periods, every cycle starts with an order-up-to-S decision followed by $R-1$ periods of zero setups. This holds for the same periods in consecutive planning runs. Thus, the sequence of respective setups is identical, resulting in a 100% *setup-oriented* planning stability (given that zero demand cannot occur with positive probability). With respect to *quantity-oriented* stability it is obvious that in the first period of a cycle a deviation between planned and executed order will emerge. On average this difference amounts to the mean absolute deviation of demand from its expectation (=forecast). For

$R = 1$ the (R, S) policy simplifies to an ordinary order-up-to-S rule just like we found for the reorder-point policies in the case of $Q_M = Q \to 0$. So the *quantity-oriented* stability level in this special (R, S) case is below 100% and is identical to the respective value for the (s, nQ) and (s, S) policy with $Q_M = Q \to 0$. For larger reorder cycles $(R > 1)$ the planning stability increases due to more and more identical non-setup periods and approaches 100% for $R \to \infty$.

A comparison of the three control policies shows that with respect to planning stability the periodic (R, S) rule is always superior to the reorder-point policies. Referring to the stability properties of (s, nQ) and (s, S) rule, a clear dominance does not exist. If we compare the stability results for varying lotsize parameters $Q = Q_M$ we can observe that the *setup-oriented* stability is higher for the (s, nQ) policy if $Q = Q_M \leq D$, but higher for the (s, S) policy if $Q = Q_M > 2 \cdot D$. In between this lot parameter interval, the relative superiority depends on the variability level of demand. When we consider *quantity-oriented* stability, analytical results for a direct comparison are only available for the range $Q = Q_M \leq D$. Here it turns out that, like for the *setup-oriented* measure, the (s, nQ) policy guarantees a higher planning stability than the (s, S) rule. From numerical investigations there is some evidence that this superiority of the (s, nQ) rule holds for the complete domain of lotsize values.

5 Long-term Stability

Long-term planning stability refers to a situation where order deviations over multiple periods of a planning horizon play a significant role in the context of nervousness. In this case, an appropriate measure of planning stability must include additional information. First, the number of planning periods has to be determined for which deviations are of relevance. Here, this number will be introduced as the stability horizon, denoted by T. Second, it must be considered how order deviations in different periods should be weighted for an overall stability metric. Subsequently, it is assumed that all periods are weighted equally.

In De Kok and Inderfurth (1997) it can be found that the analytical derivation of the stability functions for the *short-term* stability case needs a huge amount of cumbersome algebra. In the *long-term* stability context, complete sequences of T orders have to be compared and analyzed with respect to their deviations. Under these conditions the analysis and derivation of closed-form stability functions is even a lot harder, mainly because a high number of cases must be analyzed in parallel. This holds for the investigation of the (s, nQ) and (s, S) policy. The (R, S) policy, in contrast, is easy to analyze because the situation is essentially the same as under *short-term* stability considerations. So, in principle, the *short-term* stability function $\pi(R)$ carries over to the *long-term* treatment under an (R, S) rule.

Literature contributions which succeed in transferring the complex mathematical analysis for *short-term* stability to the even more complex *long-term* case are given in Heisig (1998) and Heisig (2002). There, closed-form solutions are presented for the stability functions of the (s, nQ) and (s, S) policy in the case of *setup-oriented* planning stability. For the investigation of *quantity-oriented* stability, however, no corresponding analysis has been performed so far.

Concerning *setup-oriented* stability, some results from the *short-term* stability analysis carry over to the *long-term* case. This, e.g., holds for the 100% stability property if the lotsize parameters Q and Q_M approach zero or infinity. We also find that with increasing value of Q under an (s, nQ) policy the stability metric decreases until it reaches a minimum value in the range of $D \leq Q \leq 2 \cdot D$ and rises monotonously with further increase of Q. For $T = 1$ the minimum level of stability has (by definition) the same value as in the *short-term* case and decreases slightly as the stability horizon T increases. For a large horizon and high demand variability the level of planning stability can go down to only 40%.

When we consider the stability properties of an (s, S) rule under *long-term* conditions, numerical investigations reveal some major differences to the *short-term* case. First, the stability function $\pi(Q_M)$ contains multiple points of discontinuity at which the stability level performs considerable jumps. Second, the course of the stability function between these jumps is not monotonous. However, after T jumps things change, and the stability is steadily rising with increasing lotsize parameter Q_M. Third, the point of minimum stability is not necessarily located at a lotsize value of $Q_M = D$ but can also switch to $Q_M = 2 \cdot D$ as stability horizon T increases. The level of minimum stability is somewhat lower than for the (s, nQ) policy, but different from this policy it can also rise with increasing T. Generally, a global comparison over the whole range of lotsize values reveals that both control policies perform similarly with respect to planning stability.

In Heisig (2002) the formulas for the *long-term* stability functions are exploited to gain insights into additional aspects. For example, the impact of different weighting schemes for the periods of the stability horizon is investigated. Also, the influence of forecast errors, i.e. deviations of the forecasted demand from its expected value, is checked. In Heisig and Fleischmann (2001) and Heisig (2002) it is demonstrated that the steady-state technique for analyzing planning stability can be transferred to an (s, S) control policy which is applied to a product remanufacturing context where additionally stochastic inflows of recoverable products are included.

5.1 Managerial Insights

The analytical stability research reported in this paper provides very useful insights for managing production and materials planning decisions when nervousness in the planning process plays a significant role. Next to cost and service aspects, relevant decision rules as they are given by the basic control policies of (s, nQ), (s, S), and (R, S) type can affect the level of planning stability to a very different extent. This holds for all types of stability measurement (*short/long-term* as well as *setup/quantity-oriented*). Summarizing, from the above research contributions six important managerial insights can be deduced.

1. Planning stability is only affected by the lotsize-specific policy parameter of the respective control rules. This means that the parameter choice concerning service level aspects is not relevant for planning stability considerations.

2. With respect to the stability objective, the periodic (R, S) order policy is superior to the two reorder-point policies of (s, nQ) and (s, S) type. This suggests that the (R, S) rule should be preferred as long as other aspects like its missing flexibility and its limited cost effectiveness do not argue against.

3. If *quantity-oriented* stability is relevant within the planning process, the (s, nQ) policy is superior to the (s, S) rule if *short-term* stability is considered. Since the stability level is constant for each Q value, the parameter choice for the (s, nQ) policy can be made independent of planning nervousness aspects in this case. Unfortunately, due to missing research contributions this result cannot be confirmed in the *long-term* stability context.

4. In situations where *setup-oriented* stability is relevant, there does not exist a general superiority of (s, nQ) or (s, S) rule to the other. Depending on the choice of the lotsize parameters (Q and Q_M), one policy can perform better than the other, but the differences are not serious. In general, the (s, nQ) rule with its standard lotsize may be more attractive from a logistical and organizational point of view.

5. When applying an (s, nQ) or (s, S) policy, from a planning stability point of view the respective lotsize parameter should be chosen carefully. The above analysis reveals that a parameter choice in the range between once and twice the demand forecast value can lead to considerably low stability

levels. So, these lotsize values should be avoided unless they are specifically attractive for cost reasons.

6. The level of planning stability highly depends on the variance of the random demand. Concerning the minimum stability level, unfavorable lotsize choice can lead to a stability level below 50% if demand variability is very high. This means that on average order deviations can be so large that more than 50% of worst-case differences may prevail. This holds for both setup and quantity deviations.

6 Challenges for Future Nervousness Research

In this paper, it is demonstrated that useful general insights can be provided by advanced approaches of analytical nervousness research. It would be desirable, however, if some more insights would be generated by extending this research procedure to further problem areas.

The most urgent open question refers to the analysis of the *quantity-oriented* stability of an (s, nQ) policy in the *long-term* context. It would be highly interesting if the independence of stability metric π on lotsize Q for *short-term* consideration carries over to this case. If yes, the third managerial insight in section 4 would be valid for all time-oriented instances. Only one single finding from the *short-term* analysis can easily be transferred. Since for $Q \to 0$ a basic order-up-to-S policy is valid in each period, the level of *long-term* stability is identical to the *short-term* level in this case. It is, however, an open question if this level remains constant with increasing Q values. Simulation studies of this case in Jensen (1993) and Jensen (1996) do not give a reliable answer to this question. Thus, an analytical study is needed even if it might be extremely cumbersome from a technical point of view.

A completely new research topic would be the analysis of nervousness under non-stationary stochastic demand. From a rough simulation study in Kilic and Tarim (2011) we know that the type of demand pattern has a distinct impact on planning stability under different control policies. For more general and deeper insights an analytical investigation would be needed. This might be feasible for a very simple cyclical demand scheme.

A further challenge consists in extending the analytical approach to a multi-stage production system. Simulation studies like in Jensen (1996) reveal that the *long-term* stability of the control rule at the end-item level has a major influence on the planning stability of the entire system. In this context, a lot-for-lot ordering policy at lower stages seems to be favorable in reducing the total system nervousness. It certainly is worthwhile to investigate analytically if the simulation-based findings can be generalized. At least for a simple linear two-stage system there should be a chance to extend the single-stage analysis.

In literature, there is a major debate on the effectiveness of various nervousness dampening strategies (see Atadeniz and Sridharan 2020). A widely used approach in this context is the strategy of freezing planned production orders (at end-item level) for a given number of periods. Many simulation-based investigations show that this procedure can lead to a significant reduction of system nervousness. This, however, can come at the expense of a major increase in costs and decline in customer service if the freeze length is too high. An analytical stability study for a single-level system would shed some light into the general interdependence of freezing and planning stability. Such an investigation might be feasible for a relatively simple case of a frozen horizon of one or two periods.

Another approach for nervousness reduction is proposed in the form of using safety stocks for avoiding order deviations. A respective method which can be implemented by extending the reorder-point control policies is described in Jensen (1993). Thereby, nervousness is reduced by abstaining from reorder-point triggered decisions if they deviate from former planned orders and if the inventory level is within a critical range around the reorder point. The width of this range represents an additional policy parameter. Simulation results presented in Jensen (1993) could be confirmed and generalized if an analytical study of the stability performance of such a three-parameter policy would be conducted.

The managerial insights formulated in section 4 could be enriched and extended considerably if at least some of these tasks for further research were tackled. The mathematical challenges are enormous. Against the background of potential insights, however, it would be highly desirable if ambitious researchers could be motivated to take on this job without (nervous) hesitation.

Planned Lead Times in the Age of Ubiquitous Information and Unlimited Computation

Jan C. Fransoo[2]

Abstract: *Planned lead times have effectively been central to supply chain planning, whether in conventional MRP or in Ton de Kok's Synchronized Base Stock Policy. As long as planned lead times have been around, they have also been challenged. In 2003 Ton de Kok and I formalized the reasoning around the use of planned lead times in supply chain operations planning. In this chapter, I explore whether ubiquitous access to real-time information and unlimited availability of computational capacity would eliminate the need for planned lead times.*

1 De Kok and Fransoo's 2003 Book Chapter

I am fortunate to have been a colleague of Ton de Kok at Eindhoven University of Technology since he joined the (then) School of Industrial Engineering in 1992 until my departure in 2017. We have had countless interactions on every imaginable topic. The most lasting and in-depth interactions have been those on the conceptual aspects of production and supply chain planning. This conceptual line of thinking had been a traditional stronghold of the TU Eindhoven School Industrial Engineering ever since the seminal contributions by Will Bertrand, Hans Wortmann, and Jacob Wijngaard (Bertrand and Wortmann 1981, Bertrand et al. 1990).

[2]Professor of Operations Management and Logistics, Department of Industrial Engineering and Innovation Sciences, Eindhoven University of Technology; Professor of Operations and Logistics Management, Department of Information Systems & Operations Management, Tilburg University.

When Ton de Kok and Steve Graves took up the unimaginably gigantic effort of creating a supply chain management handbook (De Kok and Graves 2003), Ton invited me to work with him on a chapter (De Kok and Fransoo 2003) for the book where we would make our conceptual thoughts explicit in the chapter's first two sections. In the remaining sections of the chapter, Ton provided a generic formulation of the supply chain operations planning problem that allowed for comparing mathematical programming formulations with those from multi-echelon inventory theory. It is fair to say that the chapter has become an evergreen, still collecting citations more than 20 years after its publication. It is arguably the only academic work that bridges the distinct perspectives of conceptual models, deterministic mathematical models, and stochastic multi-echelon models for supply chain planning.

1.1 Planned Lead Times

In our chapter (De Kok and Fransoo 2003), we argue that lead times are *exogenous* to the supply chain operations planning problem, and hence are *planned* rather than an *outcome* of the planning process. I will not repeat the entire line of arguments here, but essentially our argument is based on the fact that in supply chain planning there is a need for coordination between higher-level decision-making and lower-level decision-making, and coordination with suppliers and customers along the supply chain. The need for coordination between different hierarchical levels results from the fact that some decisions take much more time before the decision is executed in reality than other decisions.

For instance, for a complex module such as a Zeiss optical system for an ASML lithography machine, it may take multiple years between the moment of placing the order and the moment of receiving the component in Veldhoven. This implies that this ordering decision needs to be made much longer in advance than other ordering decisions for modules or components with much shorter lead times. The earlier ordering decision may then be based on a different demand forecast than a later ordering decision, resulting in coordination problems. If planned lead times were endogenous, this would cause more or less continuous changes in orders with suppliers.

Planned lead times are the foundation of almost any real-time planning system. For instance, higher-level planning decisions (with longer planned lead times) are made less frequently than lower-level planning decisions (with shorter planned lead times). Effectively, organizations that are in control of their performance have set their planned lead times, planning frequencies, and arrangements with suppliers and customers such that the supply line is balanced.

1.2 Non-stationarity and Supply Stochasticity

Arguably, the reasoning that we developed in our chapter is based on a stationary perspective of the supply chain. It requires quite a bit of advanced conceptual thinking to maintain this line of thought for supply chains where demand is growing, declining, or seasonal. The concept of planned lead times is attainable as long as supply and demand are balanced. This implies that a company facing substantial growth in demand will need to invest in its capacity and the capacity of its suppliers ahead of time such that lead times remain more or less stable. This is essentially reflected in the principles of factory physics (Hopp and Spearman 2008) that have been developed on foundational analyses in queuing theory for manufacturing (Buzacott and Shanthikumar 1993). Similar lines of reasoning can be maintained for declining and seasonal demands.

If supply is stochastic, for instance due to high capacity utilization, our line of reasoning would argue to include slack in the lead time such that a certain percentile of supplies would arrive within the planned lead time. Such slack would then be combined with higher inventories in the supply chain.

1.3 Ubiquitous Information and Unlimited Computation

With the advent of more granular information systems, enabled by decreasing costs of sensing data, researchers and consultants increasingly argue – albeit often implicitly – that waiting for planned lead times to realize is too costly, especially in stochastic systems facing higher levels of disruptions. The argument is that better and more current information on the real-time status of the supply chain has become available at low cost. Moreover, computational abilities have increased, allowing for solving models at high speed and low cost. The question then arises if the concept of planned lead times still makes sense in this time of ubiquitous information and unlimited computation. One could argue that decoupling hierarchical levels of decision-making operating at different frequencies is far more costly than simply recomputing at a high frequency. Further, one could argue that suppliers could receive signals with updated orders at high frequency leveraging the low cost and high ability of granular and frequent information exchange. In a recent interesting paper, Kasper et al. (2024) show that workload release and workload control should be more integrated, while in our reasoning these two functions should be kept separate. In the next section, I will explore whether ubiquitous information and unlimited computation imply that planned lead times can be abandoned. My conclusion is that this is not a good idea.

2 Updating Lead Times as a Thought Experiment

I now explore a thought experiment where planned lead times are endogenized or updated.

2.1 Modeling framework

First, we define L_i as the throughput time between the time of release of an order for item i and the time at which the ordered items are available for usage in other items and/or delivery to customers. If a replenishment decision $r_i(t)$ is made for item i at period t, this decision under an exogenous planned lead time L_i implies that a forecasted delivery after the planned lead time $\hat{p}_i(t, t + L_i)$ is based on the information (such as the demand forecast) at time t. New information may become available at $t + 1$, such that $\hat{p}_i(t + 1, t + L_i) \neq \hat{p}_i(t, t + L_i)$. The question then arises whether $r_i(t)$ should be updated. Note that this would be an update *in hindsight*, as the order had already been released for replenishment. Hence, updating $r_i(t)$ at time $t + 1$ would only affect $\hat{p}_i(t + 1, t + L_i)$ if L_i could be reduced by 1 period. Hence, in addition to informing the supplier about $r_i(t)$, the supplier would also need to be informed about the desired updated delivery $\hat{p}_i(t + 1, t + L_i)$.

2.2 Information and Computation

We now assume that the different firms across the supply chain that we are planning are sharing full information not only of the status but also of the planning dynamics, for instance through a detailed simulation model. If such a simulation model is continuously updated with live status information, it is typically denoted as a *digital twin*. Following the reasoning above, updating $r_i(t)$ in hindsight requires the supplier to accelerate this order through its facilities. Following queuing theory, this would imply additional capacity is needed, and the required components or raw materials would be needed earlier than planned. Additional capacity could be attained by working overtime. In many cases, however, planning overtime requires planning some time in advance. Assuming this is also covered in the digital twin's model, also that feasibility could be checked, along with checking for available component inventory. If component inventory would not be available, maybe an emergency order could be placed with the supplier's supplier, and so on.

This qualitative reasoning shows that if a perfect and complete digital twin with complete and accurate information on the detailed status of all resources were available, the feasibility of updating the plan could be checked, and the associated computations and planning decisions could be made such that $\hat{p}_i(t+1, t+L_i)$ could be attained if the digital twin were complete and information were complete and correct.

2.3 Scaling Decisions

The next part of my thought experiment is to scale this across an entire set I of items i. Each of the items has its own endogenous lead time L_i. Now matters get more complex, as somehow we need to process updates of $r_i(t)$ for every i in parallel. To check the feasibility of updating $r_i(t)$, we also need to be aware of $\hat{p}_j(t-L_i, t)$ for all components j of i. That is, we need to check the availability across J and I, which is based on the prior forecast of availability $\hat{p}_j$.

Essentially, endogenizing planned lead times requires continually updating expectations of delivery of all items across the supply chain. Note that in this planning framework, we are still only planning with expectations and not with entire probability distributions. Scaling would imply that more or less consistently we allow for

- $r_i(t)$ to be changed retroactively,
- $\hat{p}_i(t)$ to be updated accordingly,

across the entire bill of materials and the entire bill of process. That is, we do not only require the status of the digital twin to continually represent the current reality but also the complete (expected) future reality given that all decisions as executed as decided. In essence, it requires that the representation in the digital twins is identical to the complete current and (expected) future reality of the supply chain we are controlling.

Note that the future reality is the outcome of a random process that suffers from the curses of dimensionality, so the above scaling is a substantial simplification of reality. Under such random processes, Selcuk (2007), a PhD student jointly supervised by Ton and me, shows that a non-zero slack is optimal. Moreover, he also shows that frequent updating of $r_i(t)$ increases the variability of performance, and decreases long-term performance.

3 Real-world supply chains

Real-world supply chains are far more complex than what is considered in our 2003 chapter and also in the supply chain operations planning literature in general. It goes beyond the scope of this chapter to discuss these complexities in full, but I highlight a few of these to demonstrate that having planned lead times more likely allows for coordination in complex supply chains than endogenized lead times would.

First, supply chains consist of multiple companies that each are part of different supply chains. This leads to competing interests and – consequently – strategic behavior. Strategic behavior has been extensively studied in many supply chain settings leveraging a game-theoretic modeling framework. The thought experiment outlined above does not accommodate such strategic behavior. Making use of planned lead times enables firms to make contracts with suppliers and customers with such planned lead times at the core of the agreement. This is also what we commonly see in contracting, where lead times, along with volume agreements, are part of the operational alignment.

Second, in order for resources to be used efficiently, performance needs to be predictable. Planned lead times are the basis for being in control of a supply chain's performance. It may seem counterintuitive since the use of planned lead times requires the explicit inclusion of slack in the supply chain. However, well-positioned slack allows for efficient use of resources at many other places in the supply chain (Selcuk 2007, Galbraith 1974).

Third, in the foreseeable future, most supply chain decision-making will involve humans. Planned lead times do not only serve as a prediction, but also as a (behavioral) target. This implies that if information shows that a lead time may not be attained, a planner will discuss this with a firm's supplier and explore whether acceleration is feasible. This behavioral control may be an explanation for why automatic lead time updating has not been very successful in practice.

4 Outlook

As Ton de Kok retires this year, it is timely to reflect on whether planned lead times will remain at the core of most supply chain operations planning mechanisms, whether archaic such as MRP, or advanced such as De Kok's Synchronized Base Stock Control. I dare to predict that at least until the planned lead time of my formal retirement age in 2033 we will effectively see supply chains being controlled with some form of planned lead times. The good news of this is that we then know how to be in control. Given the massive uncertainties that the world's supply chains are facing, this is probably a more reassuring thought than the theoretical question of how we can eventually improve performance with truly ubiquitous data and unlimited computation.

My favorite mathematician

Vincent Wiers

1 Mathematicians and me

Since my PhD project at the Eindhoven University of Eindhoven, department of Industrial Engineering, which started in 1993, I have been doing research in the area of the human factor in production control. The idea of focusing on the human factor in planning and scheduling problems originated from the observation that there are many theories and techniques described in academic journals that are nevertheless not used in practice. Most techniques to generate plans and schedules were developed by academics working in Operations Research, which can be regarded as applied mathematics. So why are most of these techniques not being applied? Human planners and schedulers might be doing some irrational and weird things that make it difficult to implement the beautiful theories. If only these human planners would realize how academics would be able to help them!

31 years later, these thoughts still make me smile. The stuff that I have seen since then in planning departments has made me realize that academics who think along these lines live on another planet. Even happily so. Practice can be scary, messy, and confusing. Many never seem to leave the sterile and safe surroundings of the university buildings. Or perhaps they do have some students doing some projects in companies, but they are never part of a real project themselves, where a process changes or a planning system needs to be implemented, with a deadline and limited budget. With users that don't take for granted what you are offering them. Users that will make it clear that there is so much to learn, before you

can even grasp the issues they are dealing with. By the way, there are armies of consultants that have similar blind spots.

Many of the articles I (co-)wrote start by introducing the gap between theory and practice, followed by a description of how human planners solve real-world problems. After a few months into my PhD I would skip all papers that did not contain an application in real-life. There were simply too many of them – papers presenting techniques to solve baby problems. When a test would be there for the technique, it would be a computer simulation. Babies in kindergartens.

In 1995, the number of papers on scheduling techniques was estimated to be more than 20,000 (Dessouky et al. 1995) and almost all of it was destined to end up in the catacombs of scientific literature. Or at least, that was my view. Still, I agree with the following quote from the mathematician I am writing this piece for: "there is nothing more practical than a good theory." If only we could develop good theories for planning and scheduling processes. Theories that could be applied to real problems. Developed by academics that also understand practical problems.

On the theory-to-practice scale, I would see mathematicians as the ones closest to theory – the kind of academics who oversimplified things. Most mathematicians who were doing research in the field of production planning and scheduling were focusing on problems that were simplified beyond recognition: one or two machines, known arrival times of jobs, and 10 jobs or so. Such baby-problems are so far removed from practice that according to my academic brother-in-arms Ken McKay "perhaps a different name for the research should be considered." (McKay et al. 1988)Together with Ken, I have written two books and a number of articles on planning and scheduling problems. There was not much math in these books and not a single Greek character.

So how likely would it be for me to write a book together with a mathematician? As Will Bertrand told me at an event in 2018: "When you would have asked me, years ago: who are the most unlikely couple to write a book together? I might have answered: Ton de Kok and Vincent Wiers". And it was precisely the two of us who were presenting our book (Wiers and De Kok 2017) at that event, with Will being our guest of honor. What happened?

2 Rejoining TU/e

After I obtained my PhD in 1997, I left the university to implement decision support systems for planning in practice. At the same time, I was looking for a part-time position at the university to combine research and my consulting work.

This turned out harder than I hoped for as universities were offering less of such positions. In the year 2000, I was having a dinner to celebrate the PhD of a friend and was sitting next to Ton de Kok. I told him about my efforts to rejoin the university and that I had not succeeded so far. He replied by saying, perhaps you have been talking to the wrong people. Shortly after, I rejoined the university with a part-time contract, thanks to Ton.

It turned out that both Ton and I have a passion for developing planning systems that work. The industry standard at that time was called Material Requirements Planning (MRP), succeeded by Manufacturing Resources Planning (MRP II). In my professional work, the planning systems that were developed had to work together with MRP. It was the time when some of us still referred to planning decision support systems as Finite Capacity Planning (FCP). The background of this was that they would work in conjunction with MRP systems – where MRP would organize the material coordination, FCP would do the capacity planning based on the MRP output. At that time, I was assuming that MRP and FCP were complementary – MRP doing material and FCP doing capacity planning. The then new term for FCP systems was APS – Advanced Planning & Scheduling. By the way, there is an even newer term: Supply Chain Planning & Optimization (SCP&O) but Ton and I refuse to use it.

I learned that one of Ton's passions was to find solutions to stochastic planning problems. He had designed a technique, implemented in a decision support system, to set the inventory level for items on multiple levels in a high-tech supply chain. The system had actually been used by planners and brought significant benefits for the company. In 2004, I heard about Ton taking part in the Edelman award competition, with the project titled "Collaborative Planning at Phillips Semiconductors and Phillips Optical Storage" described in De Kok et al. (2005). The tool that Ton had designed yielded quite impressive results and it could be seen as an improvement of the decisions that would normally be taken by an MRP system. Thanks to Ton, I started to realize that MRP is flawed not only in the area of capacity planning but also in material and inventory planning. Furthermore, he also showed that planning techniques should take uncertainty into account, instead of treating information about the status of the supply chain as deterministic facts.

Roughly between the year 2000 and 2010, I spent time at the university to do joint research and write publications, mostly with Jan Fransoo, a colleague professor in the same department as Ton's. Around 2005 I was also asked to support the European Supply Chain Forum (ESCF), which is a collaboration between a group of companies and academics, mainly from TUE. Ton was one of the founders of this group and for a long time the scientific director. The ESCF turned out to be a great success in bringing companies and researchers together, creating opportunities to do joint projects, exchanging students, and deepening insights in how theory can impact practice and the other way round. In other words, another example of Ton building a bridge between academia and commercial companies. During that time, I literally saw Ton as the driving factor behind the ESCF. It is just one of the initiatives that I have seen him take, putting all of his energy into it and aiming for high rewards. The list of companies that are participating is impressive, and the ESCF has created enormous value for both the university and the member companies.

Because of the cooperation with ESCF, I got to know Ton better and we had long discussions on planning hierarchies, MRP (Material Requirement Planning) versus APS (Advanced Planning & Scheduling), and how to design decision support. Ton turned out to be able to look beyond the rigor of mathematical models. As I described already, I did have a predisposition about mathematicians, but Ton turned out to be different. On the one hand, he had been developing models that were able to deal with the stochastic nature of real-world phenomena like customer demand. He was able to explain to me how my decision support systems were flawed, in not being able to deal with uncertainty – other than replanning. Also, he was open to discussions where I would openly question the applicability of mathematical models and standard production control frameworks.

Ton is one of the few academics who combines his academic approach with common sense and asks himself the question: Would this work in practice? Many academics would avoid having discussions with me as I would only spoil their party by declaring the baby problems as completely unrealistic. However, Ton was curious to hear about experiences in practice, and I was interested to hear more about the errors in the assumptions I was making all the time. Ton taught me about the shortcomings of MRP, which were much more extensive than I knew. Most of all, we had fun and each discussion brought new ideas and insights.

3 Course on Advanced Planning & Scheduling

In the year of 2012, Ton and I kicked off the course on Advanced Planning & Scheduling, which was back then a course for PDEng students. Quite naturally, we divide the tasks along our expertise areas: Ton focuses on some of the quantitative modeling aspects of supply chain management, and me focuses on the implementation aspects of APSses and the human factor. The course description was as follows.

"Improving supply chain planning and scheduling often involves the use of Advanced Planning & Scheduling systems. The design and implementation of such systems requires an APS specific approach, as many projects are not (completely) successful in practice. When applied correctly, APS systems can generate large operational savings (such as the case for Ton's Philips case) and improve service levels at the same time. The APS needs to fit the human planner/scheduler, the physical system to be controlled and the organizational context. Attention will be given to the characteristics of various types of APSses, the structure of models, how to apply optimization and how to make sure the projects are carried out successfully."

For the quantitative modeling of supply chains for the purpose to design APSses, Ton largely draws from the models as described in the handbook (De Kok and Fransoo 2003). The assignments that Ton offered the students would force them to study the MRP logic in detail and to analyze the differences between different techniques for inventory planning. I remember that especially in the first years of teaching our course, I had difficulties trying to understand the material that Ton was presenting. Some students would also point out the different viewpoints that Ton and I were taking. This exam question from the year 2013 nicely illustrates how the different worlds compared:

Reality is stochastic, as we cannot know the future demand or the exact timing of future deliveries of items to stockpoints. Unfortunately, stochastic models are intractable for realistic supply chain models. Therefore, APS engines are built on rolling schedule concepts and mathematical programming models. Explain how these models take into account uncertainty. Explain how this approach is fundamentally different from the way single-item single-echelon inventory models like (s,S)-models deal take into account uncertainty.

Thanks to Ton, I started to realize that all my work in real-life was based on the assumption that information was perfect, although all the time new information would come it that superseded the old. And that we were dealing with this by using a rolling planning horizon and putting some buffers here and there – like leadtime buffers, overcapacity, inventory, delivery time buffers. When we discussed these differences with students, I would admit that I have no other option – as far as we know, there are no models that can be used by optimization techniques, which include material and capacity constraints, and that can take uncertainty into account.

4 Some of our favorite discussion topics

4.1 Goodsflow Control, Production Unit Control and Leadtimes

Where Ton and I have had fierce discussions about and at some point 'agreed to disagree', is the topic of goods flow control versus production unit control, and the idea of having fixed lead times as a means of connecting the two. To clarify this topic, I will need to explain how the 'Eindhoven school' teaches these concepts.

Production systems are often decomposed into a hierarchically organized planning and control structure to reduce complexity. This approach is also known as Hierarchical Production Planning (HPP). Bertrand et al. (1990) distinguish between goodsflow control, which concerns planning and control decisions on the factory level, and production unit control, which concerns planning and control decisions on the production unit level. The goodsflow control level also coordinates the up- and downstream production units. MRP can be seen as a concept that operates at the goodsflow control level.

A production unit is an outlined part of the production process of a company: it produces a specific set of products from a specific set of materials or components, with a specific set of capacity resources. Depending on the complexity of a production system, a production unit can be a single machine, a hall of machines with personnel, or an entire factory. Scheduling can now be defined by the organizational function that deals with planning and control decisions at the production unit level. This is depicted in the following picture:

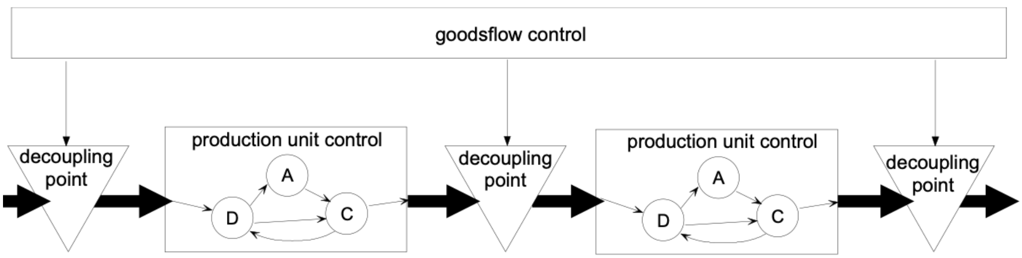

Figure 10: Goodsflow control and Production Unit Control according to Bertrand et al, 1990

Goodsflow control is normally done by MRP. In a Make-to-Stock (MTS) environment, MRP will be triggered by signals that end item stocks need to be replenished. And in a Make-to-Order (MTS) environment, MRP will be triggered by a customer order. APS systems can also do goodsflow control, tackling some of the material coordination issues that MRP suffers from. The production unit control is normally referred to as scheduling, and this is the most detailed control level before production takes place.

I have told Ton several times that I believe that this key building block of production control, as taught to many students in Eindhoven, is wrong. The reason being is that I have never seen a scheduler, that is, someone who controls the production unit, who does not check material availability – both on the input and the output side - while the framework simply assumes that goodsflow control takes care of materials and in scheduling you can safely assume that the ‘released’ orders can be started. Furthermore, the problem of material allocation can be quite complex in practice, for example, when discrete pieces of material with certain dimensions need to be allocated to downstream planned orders to fulfill a specific quantity. This distinction between material balancing on the goodsflow level - where only the quality balance is checked – and the detailed allocation – where the specific material needs to be selected – seems absent in literature. MRP seems to assume that the actual material *allocation* is a trivial decision, once the material balancing has been done on the GFC level. For many practical situations, this is a very wrong assumption. It can even be that quantity wise, the material balance seems in order, but when assigning specific pieces of material, there is a shortage. The standard GFC/PU control framework seems to completely ignore this issue.

4.2 Order Release

The term ‘order release’ has caused quite some confusion on my side and discussions between Ton and me. Ton uses this concept often, and whenever he did that, I would struggle with the meaning of ‘release’ – perhaps because I am obsessed with using names for terms that say something about the meaning of the term. Release – would that mean that the order does exist before, but it is made visible to the scheduling level at the moment of release? However, this seems not correct, as an order can apparently be released at the moment of creating the order. So, in my mind, I replace the word ‘release’ with ‘planning’ – and now I more or less understand what Ton is talking about.

4.3 Leadtimes

In goodsflow control, it is normally assumed that the leadtime for an item is constant. I have always fiercely challenged this assumption, as the leadtime will vary according to the production unit’s workload, the routing chosen, and nasty things that can happen on shopfloors, such as quality issues. As for the workload part, the Eindhoven school would offer workload control as the mechanism to keep the leadtime constant. However, that would mean that there is a set of orders that is in the production unit and a set which is not. When there is too much work to be done, there will be many orders in the same period, competing to be 'released' and goodsflow control does not have an answer to that.

But sometimes a discussion must sink in some time before you realize how the different viewpoints relate to each other. And the longer I thought about it, I realized that Ton had a point. Perhaps leadtimes are not always constant, but there are good reasons to state that they should be, and that organizations should put some effort into achieving that. When decomposing production control into levels, you need to be able to make assumptions on what the lower levels will do without knowing all the details. Constant or predictable lead times - whenever possible - contribute to a healthy planning practice.

4.4 MRP drawbacks

Ton has really opened my eyes to what MRP can and cannot do. As I explained earlier, I saw MRP as the technique that can do material coordination, and APS systems that can do capacity coordination. However, one of the assignments that Ton handed to the students clearly showed that MRP is actually seriously flawed in material coordination. First, MRP does not synchronize supplies – meaning that when an item needs two components and one of them is postponed, the other one is still ordered to be produced in time. Second, the MRP generated plan can contain infeasibilities such as negative stocks for upstream items. Hence, MRP is not very suitable to act as a planning technique on the goodsflow control level.

5 The end of 1CM150

The retirement of Ton also marks the end of the course on Advanced Planning and Scheduling. I have mixed feelings about this and I am sure Ton has them as well. On the one hand, it was fun to teach students about applying theory and techniques to practice. Every year, the students would write down in the course evaluations that they valued the relevance of the course for real-life and the war stories that Ton and I would tell them. On the other hand, it would be hard to replace Ton for the course, and it also feels logical to conclude 13 good years of teaching the course. I am sure that Ton and I will continue to have discussions on APSses and we already have an idea that we want to follow up. Doing fun stuff with Ton on production control topics will never end!

An Appreciation of Professor Ton de Kok

Stephen C. Graves[3]

1 Introduction

Professor Ton de Kok is a wonderful friend and long-time colleague with an exceptional record of impactful research. I am delighted and honored to be able to participate in this event, celebrating his career and its accomplishments.

Ton is an outstanding scholar, best known for his research on the modeling and optimization of complex inventory systems that arise in virtually all supply chains of any interest and importance. He has a distinctive background, combining fundamental training in operations research with full-time practical experience in analytical problem solving at Philips. This background is well reflected in his problem-driven research. Since returning to academia, Ton has interacted extensively with industry to both uncover relevant research opportunities and apply his research to improve practice. He has a long record of addressing important real-world operational problems in supply chains and doing so with novel modeling, along with rigorous methods and analyses. His research contributions entail better methods and better decisions for a host of inventory-related planning problems that arise in production, distribution, and maintenance settings.

One outstanding example is his development of Synchronized Base Stock policies. These policies provide an applicable way to model and analyze the inventories in supply chains with complex product structures, with planned lead times, and with

[3]Abraham J. Siegel Professor of Management, Professor Post Tenure, MIT Sloan School of Management, Cambridge MA, USA. sgraves@mit.edu

stochastic demand. These policies encompass many simpler policies and provide a unifying framework for the analysis of these systems. They also lend themselves to cost optimization accounting for stock-outs, and the consideration of capacity constraints, in order to get the right inventories in the right places throughout a supply chain. Ton has demonstrated the effectiveness of these policies in a number of industrial projects, including applications at P&G, ASML, Philips, Hilti, and Dow, to name a few (De Kok 2018b).

Ton has also done noteworthy work in production planning, particularly hierarchical planning, as applied to supply chain planning, and in the control and scheduling of production/inventory systems. In each instance, the work again reflects his distinctive abilities to marry practical considerations and needs with methodological and modeling innovations. Some other research highlights in which he has addressed real problems with impact, include (i) research on vehicle routing with stochastic travel times, which arose from a project with a transportation company and a dairy company (Taş et al. 2013); (ii) a multi-item, multi-echelon model for setting safety times in make-to-order production systems, in collaboration with ASML (Atan et al. 2016); and (iii) a stochastic multi-echelon model for facilitating collaborative planning at Philips Semiconductors, to 'tame' the bullwhip. The application of this research resulted in substantial financial benefits at Philips Semiconductors (De Kok et al. 2005).

2 An Overview of Research on Multi-item, Multi-echelon Systems

I will focus the remainder of this note on Ton's 2003 publication (with Jan Fransoo) in the handbook that we co-edited (De Kok and Fransoo 2003). I chose this because it captures much of Ton's work and thinking up to that time on multi-item, multi-echelon systems.[4] I will first provide a high-level recap and critique of the contributions in the chapter; and then I will follow this with some comments and observations, with special attention to opportunities for future research. In these comments, I will try to incorporate updates on Ton's more recent research on this general topic area.

The focus of the chapter is on Supply Chain Operations Planning (SCOP), which is a subproblem (or subfunction) within Supply Chain Management. Supply Chain Operations Planning entails the determination of the release (or scheduling) of materials and resources across a supply chain, with an objective of obtaining

[4] All of my positive comments about the publication also apply to the co-author Jan; but as this piece is an appreciation of Ton, I will just refer to Ton to keep the text simple.

customer service targets at minimum cost. The chapter formulates the problem with both novel modeling and an emphasis on applicability; it then presents two viable solution approaches, and contrasts and compares them.

The SCOP problem is framed in the spirit of hierarchical planning and is intended to address medium-term (or short-term) operational decisions. One assumes that long-term decisions have been made and are given, related to the design and sizing of the supply chain. These decisions involve the network design and capacity planning for the supply chain. The resulting supply chain can then be modeled as a network of transformation activities, where each activity converts a set of inputs into an output. Activities include production, assembly, transportation and procurement. As a simplification, there are no by-products, so each output is a single item at a location. Outputs can be inputs to other activities, or be end products that serve customer demand (or both). Each activity can have resource requirements, in order to activate the activity. And each resource may be shared across multiple activities. Finally, for each activity, there is a pre-determined planned lead time: the amount of time between the initiation of an activity and its completion. It is presumed that these planned lead times are based on the capacity planning that has been done at an upper level of the planning hierarchy.

The SCOP problem is given a demand forecast (and/or a demand characterization) for the end products of a supply chain, and then determines the material release for each intermediate item and end product; the material release entails deciding whether to initiate a replenishment or not, and for how much. Implicit in this decision are other complementary decisions to assure that the material release is feasible. For one, the plan needs to assure that there are sufficient resources to accomplish the planned material releases within the planned lead times. This may require decisions on how to allocate a resource across multiple activities as well as across time. In some contexts, it may be possible to make short-term adjustments to a capacity level, which would be another resource decision to include. In addition, there may be material allocation decisions for items that are inputs to multiple end products (or intermediate items).

Underlying SCOP is a planning horizon, divided up into time units (e.g., days or weeks). Whereas the plan may lay out these decisions for the entire horizon, only the decisions for the immediate time period are typically implemented, and then the planning process is repeated in the next time period, in a rolling horizon fashion.

Ton then develops two approaches to solving the SCOP. Key to each of these approaches is a very insightful assumption: “any SCOP concept can only release items for usage in a transformation activity if it is physically available at the moment of release…" This says, for instance, that we should not place a replenishment or production order if we know it is infeasible, either due to limited production resources or due to limited material availability. And as such, this seems a quite reasonable assumption. A consequence of this assumption is that there will be no backlogging of items that only experience dependent demand. However, as Ton notes, the traditional planning logic in MRP (material requirement planning) will violate this 'no backlogging' assumption, and will generate requirements that cannot be fulfilled due to material unavailability, and hence result in 'planned' backorders.

The first solution approach is to formulate a linear program that can capture all of the complexity of the network of activities, as well as the resource constraints and the no backlogging assumption. Some novelty is needed to also model the planned lead times. The input to the linear program is a demand forecast over the planning horizon, and there is no accounting for demand uncertainty. However, one can find the safety stocks for the end products to achieve some desired end-product service levels, based on a very interesting, novel observation. One can add end-product safety stocks to the linear program, and then, with an adjustment to the objective function, show that the solutions to the linear program are invariant to the safety stock level. Then from a single (long) simulation of the linear program, one can determine how much safety stock is needed to meet the service level for each end product.

The second solution approach is based on the methods analyzing multi-item, multi-echelon systems under stochastic demand. Here, the multi-echelon theory is yet to accommodate resource constraints, beyond relatively simple systems; so, the development in the chapter assumes no resource constraints.

The existing theory provides optimal solutions for the serial systems, and for pure assembly systems (i.e., one end product), and then very good solutions for divergent systems. Ton extends this literature and adapts it to the more general networks as exist in complex multi-item supply chains. As assumed in the literature, he starts with echelon base-stock policies, but these need to be adapted in a couple of ways. One is to abide by the no backlogging assumption, which entails synchronizing an item's orders with those of other items that share end-products and that have longer lead times. (i.e., the longer lead-time items impose a constraint.) The second adaptation is the need for allocation policies when inventory is short.

The building block for solving SCOP for general networks is the analysis of divergent systems. This analysis has traditionally relied on a so-called balance assumption, which assumes that any desired allocation is feasible. For divergent systems, Ton had previously characterized the optimal solution, under the balance assumption, and found that the determinations of the base stocks are analogous to generalized newsvendor problems (Diks and De Kok 1998) . However, the optimal solution is hard to implement because of the complexity of the allocation rule. Ton finds, though, that a simpler, linear allocation rule works very well, and that the base-stock levels and allocation fractions can be computed very efficiently for divergent systems (Diks and De Kok 1999).

The proposed solution for general networks is developed from the analysis of divergent systems. There are two key challenges in modeling general networks. One is that order synchronization is complicated by the absence of a clear hierarchy in decision making about order releases. The second challenge is deciding how to allocate an item's inventory across successors, when there is insufficient inventory. To overcome these challenges, Ton shows how to decompose a general network into a set of divergent systems, which can then be solved (De Kok and Visschers 1999). The decision hierarchy is given by the design of each divergent system; inventory allocation is now based on linear rules, as shown previously effective but now modified due to the more complex network structure. The devised solution is named a Synchronized Base Stock (SBS) policy and is quite creative, enabling the synchronization of order release decisions over time for general networks.

The chapter provides a numerical comparison of the two solution approaches for a two-echelon test case with four end products and seven components, and no resource constraints. The solutions were compared by simulation with randomly generated demand. The SBS solution performed much better than the linear program solution, requiring 10 to 15% less inventory investment. It seemed that the linear program did a poor job of rationing shortages, possibly due to solution degeneracy. Furthermore, the SBS solution was quite robust over a range of settings for the test case.

3 What to Like?

There are many positives in this publication, some related to the formulation (or modeling) of SCOP, and some related to the proposed solution methods. Some of the highlights that come to mind are the following:

The modeling of SCOP is an important contribution.

- The framing of SCOP within a decision hierarchy, as being medium-term planning. As such, the formulation assumes that long-range capacity planning and network investment decisions are made. The SCOP then addresses how to place and deploy inventory to meet customer service targets over a short to medium-term horizon. This framework is important to ensure the relevance for the research as a recognition of how organizations do (or should do) material and resource planning.

- The introduction of order synchronization as a required (or desired) feature of an operational plan, by which all planned releases need to be consistent and account for the material availability constraints imposed by the longer lead-time components. This is a distinctive feature of the model that adds to its value and applicability.

- The presumption of planned lead times is also an important feature for similar reasons: this increases the applicability of the research, as most planning systems rely on planned lead times in one way or another.

Both solution methods have several positives:

- They both apply to general networks, with few limitations. And both methods account for order synchronization, and both are quite feasible computationally.

- The linear program accounts for resources, and explicitly models the planned lead times. The linear program can determine the end-product safety stocks needed to achieve service-level targets.
- The SBS method explicitly accounts for demand uncertainty, and allows safety stock anywhere in the network, as needed.
- The structure of the SBS policy is particularly appealing, being pragmatic, robust and interpretable, and with the echelon base stocks being set with the newsvendor fractile. This significantly facilitates support for the general use of this kind of inventory policy.

- The ongoing efforts to validate these methods (particularly SBS) through industrial projects is especially noteworthy and critically important; this not only establishes the value from the research but also provides the feedback for future development and improvement of the solution methods (De Kok 2015),(De Kok 2018b).

4 What to do Next?

In closing, I will provide some thoughts for future research, which reflect some of my interests and will include several gratuitous references to my papers. I also note that some of my suggested questions may have already been addressed in the literature, and that I may have overlooked some of Ton's recent research that is relevant to these questions. Nevertheless, with these qualifications, here is what comes to my mind (in no order):

- How to set planned lead times? As noted, SCOP assumes that planned lead times are given, which is a quite reasonable and pragmatic assumption. But it begs the question of how these should be set, and whether planned lead times are an operational or a tactical decision? Whereas we have developed extensive models and methods for setting inventories, there is very little systematic knowledge on how to specify these critical planning parameters. One exception is Ton's work with AMSL in which they model the planned lead times for a module assembly system and then develop and apply an algorithm for setting them (Atan et al. 2016). In Graves (2022), I discuss the planned lead times as a tactical decision for balancing resource requirements and cycle time, and provide a few ideas and even more questions.

- How to consider different review periods for different items? In different parts of a supply chain, there are likely different review periods, and/or different economies of scale for replenishment activities. As such, it may not be realistic to assume a common review period. For general networks, I think this is still an open question that seems important. Possibly one can extend our multi-echelon, multi-item models to accommodate different review periods; alternately, one might envision a coupled pair of interacting models, one that determines how to handle demand uncertainty (e.g., safety stocks) and another that determines how to exploit economies of scale (e.g., review periods or lot sizes).

- How to coordinate across a long supply chain? In long, multi-organization supply chains, there can be competing and conflicting interests which may make it difficult (or impossible) to implement an optimal operating strategy (even within a single company). One challenge is whether to decouple the chain into smaller supply chains that operate independently, and if so, how to decide and accomplish the decoupling points. Another challenge can be how to construct contracts that coordinate across the different parties and result in a high-performing supply chain. In Schoenmeyr and Graves (2022), we provide one attempt at these questions, as a possible starting point.

- How to account for forecast evolution? The planning for many supply chains is driven by a dynamic forecast process. In each period, this process generates a vector of demand forecasts, with a demand forecast for each period in the forecast horizon; in the next period, the forecast vector is advanced one period and revised based on the demand realization in the current period as well as on other market information. Typically, the accuracy of each particular forecast will improve with each update, i.e., the 3-week ahead forecast should be more accurate than the 4-week ahead forecast. But the rate at which forecasts improve with each update is very context dependent. And, most multi-echelon inventory models do not capture these dynamics (an exception is Schoenmeyr and Graves (2009); but there is much room for improvement.)

- How to account for other countermeasures, including human intervention? There should be more recognition that inventory is not the only tool for handling uncertainties. In Graves and Willems (2000), we observe that when faced with excess demand, "... a manager resorts to other tactics to handle the excess demand. A manager might use expediting, subcontracting, premium freight transportation, and/or overtime to accommodate the windfall of demand." We go on to argue that the planning of safety stocks should somehow account for these other countermeasures. More recently, Ton has undertaken some very interesting research on modeling and measuring such countermeasures, which he terms human interventions (De Kok 2018a); I hope we will see more efforts on this question.

- Can we find even better allocation rules? The multi-echelon literature for divergent systems relies extensively on allocation rules derived using the so-called balance assumption. There is substantial numerical work showing that the resulting control policies generally work very well, even when the balance assumption (which equates to cost-free transshipment) does not hold. Nevertheless, one might ask if there are allocation rules that are always feasible and will work well. For instance, Graves (1996) proposes for divergent systems a naive allocation rule that is feasible and tractable, but clearly non-optimal; I suspect there are opportunities to devise better allocation rules that are feasible, practical and perform well.

A Tribute to Ton de Kok

Hau Lee[5]

It is an honor for me to recognize the contributions and the professional relationships that I have had with Professor Ton de Kok. Ton has been a pillar in the production and operations management area, and a leader of our field.

My first encounter with Ton was on my first trip to Eindhoven University, during which I also brought with me Dr. Corey Billington, a former student of mine who was then a manager at Hewlett-Packard Company. On this visit, Ton, together with Will Bertrand and Jan Fransoo discussed possible avenues that Eindhoven, Stanford and HP could pursue together. We all emphasized excellence in teaching and research, but we also recognized the importance of working with industry as close partnerships. That planted the seed of a long journey of collaboration.

Shortly after our visit, in 1996, Ton and Jan formed the European Supply Chain Forum, housed at Eindhoven University. They were gracious in inviting me to be an advisor to this Forum, and I have benefited tremendously from my association with the European Forum. We immediately established a sister relationship for the European Supply Chain Forum and the Stanford Supply Chain Forum. This sister relationship creates opportunities for members of both forums to attend each other's events, share research outputs and newsletters, and promote faculty exchanges and student collaborations.

The two supply chain forums have been strong supporters of one another. We have shared our ideas on themes for workshops, and tried to help promote each other as much as possible. I am grateful to the leadership of Ton in this regard.

[5]Stanford University

I still remember fondly that, at one of the European Supply Chain Forum roundtables where I attended, Ton arranged for the group to watch a PSV game in the evening. It was my first football game in the Netherlands, and it was fun.

In 2000, we recognized the importance of Asia as a major part of the global supply chain. Together with the Supply Chain and Logistics Forum at the Hong Kong University of Science and Technology, the European Supply Chain Forum and the Stanford Supply Chain Forum organized a workshop on Asian Supply Chain Challenges and Opportunities, as our first three-way collaboration. The workshop was held in Hong Kong, with participants coming from industrial members of all three forums. This event has helped us understand much better the realities of supply chain management in this part of the world, as well as provided us with ideas for research. In 2004, the three institutions collaborated again, and this time, we included the China European International Business School in Shanghai, on a conference on "China at the Crossroad of the Global Supply Chain." A record 200 industrial participants in attendance, spirited exchanges on China's supply chain issues, and site visits to factories and logistics hubs, marked the success of this event, held in Pudong, China. This was followed by another global conference in 2007 that the three forums jointly sponsored, hosted by Zhejiang University in Hangzhou. Below was a picture taken at the Hangzhou conference.

Stanford and Eindhoven were considered to be the pioneers in starting something called Supply Chain Thought Leaders Roundtable. We felt that supply chain research was emerging everywhere, and there was a need for the top researchers to get together to share, exchange, and identify new frameworks and paradigms. Stanford hosted the first one in Santa Cruz in 1998, and we were able to assemble 30 of the world's top academic and industrial leaders in supply chain research at this roundtable. Together with Jan Fransoo, Ton hosted the second Supply Chain Thought Leaders Roundtable in Corsendonk, Belgium. In 2010, Ton and Jan also organized the Roundtable in Breda and Rotterdam. Enclosed is a picture taken there. The Roundtable is now a new tradition for the top researchers in supply chain management, and Eindhoven and Stanford are proud to be the "founders" of this event. Here is a picture of Ton and Jan at the Roundtable of 2007 in Kobe, Japan.

The primary focus of Ton de Kok's research has been on stochastic production and inventory models. His modeling work has been motivated by real-life systems, such as those at Philips Electronics. He has made contributions to queueing theory, inventory theory, supply chain management, and stochastic vehicle routing.

Multi-item, multi-echelon inventory systems have been a major subject for many researchers in operations research. Ton has extended and generalized the early

pioneering work of Clark and Scarf (1960) and Eppen and Schrage (1981) for complex structures such as divergent and assembly systems. He was most recognized by the development of Synchronized Base Stock (SBS) policies for the control of complex inventory systems. He also came up with a recursive method so that the SBS policies can be solved with linear complexity to ensure scalability of the approach. Over three decades, he has extensively validated the results under SBS policies in real-life situations through projects at P&G, ASML, Philips, Hilti, and Dow, etc.

A major result of his inventory research was on how optimal installation stock policies with a fixed reorder level or safety stock satisfied the Newsvendor fractile. This provided support for the general use of this kind of inventory policies.

Over the years, Ton has helped to implement his inventory research results at Philips Electronics. Together with his graduate students, he also expanded such implementations at other companies, bringing great financial benefits in the order of hundreds of millions of euros to these companies. To create more long-lasting relationship between academic research and industry practice, he played a major

leadership role to build up the Parts Business Forum (PBF). Since 1993, he helped to expand PBF from having five companies as members to 60+ now. PBF has become the Service Logistics Forum. Ton invited colleagues from other Dutch universities to participate, creating a strong Dutch research community, which later became one of the pillars of Dinalog, the Dutch national knowledge institute on Logistics and Supply Chain Management, established in 2010.

As noted above, Ton is a research pioneer who has made tremendous theoretical contributions in multi-echelon inventory systems. His work was also grounded well in industry applicability. I will not dive into the academic contributions that he made here. At the same time, he recognized the importance of globalization. As evidenced by his active championing of cultivating shared understanding and exchanges in industrial forums held in Europe, Asia and the US, his work on supply chains has global significance.

I personally have gained tremendously from our collaboration. Thank you, Ton, for your admirable achievements. May I take the final moment to congratulate you on your new phase of career and life in retirement. I will always cherish our friendship.

Workload-Dependent Lead Times in Production-Inventory Systems

Reha Uzsoy

1 Introduction

After 35 years in academia, it feels like Ton has always been on my radar as a leading academic in the domain of production-inventory systems. During the first phase of my career, when I worked primarily in deterministic scheduling, our interests with Ton did not overlap a great deal. My closer interactions with Ton began with an INFORMS meeting in the early 2000s, where a group of us including Jan Fransoo (then at Eindhoven), Ton, and Karl Kempf of Intel went to dinner and had a very stimulating conversation. Ton and I began to correspond, I hosted several (very impressive!) Eindhoven students at North Carolina State, Ton came to visit us, I attended workshops at Eindhoven and visited Ton there, and Ton and I served on each others' students' doctoral committees. Around the same time Ton took his latest position with CWI, I took a temporary position as a Program Officer with the National Science Foundation, which allowed us to share views and experiences on the management and organization of research activities. Although we have only recently begun to work formally on a specific project together, I have had the pleasure and privilege of learning from a true master of his trade.

Several interesting characteristics distinguish Ton from many fellow academics. Although he is deeply knowledgeable in, and deeply committed to, rigorous mathematical analysis, he does not let the mathematics obstruct a clear idea of the practical problem at hand, with the relationship between the mathematical models and the practical problem always kept clearly in view. His constant involvement with industry through students, funded projects and sabbaticals, is truly a cause of envy to those of us in the U.S. where such relationships are difficult to sustain over long periods of time. It is also striking that Ton has made significant contributions to both stochastic inventory theory and deterministic production planning, areas requiring very different perspectives and mathematical tools, and hence are often difficult for one person to master.

Finally, I cannot pass without noting Ton's impact on those around him. He invariably brings out the best in those around him, with helpful and thoughtful comments and suggestions. For many of us, the image of Ton at a conference, with his "chrome dome" six inches above most of the crowd and his characteristic grin, engaged in animated discussion with a group of colleagues and thoroughly enjoying himself, will remain a treasured memory.

2 Motivation

While the greater part of Ton's work lies in the domain of stochastic inventory models, with the development of effective computational procedures for general production-inventory network structures as a clear theme, the handbook chapter with Fransoo (De Kok and Fransoo 2003) is a very important contribution in its own right, very much in the Eindhoven tradition of developing frameworks for problem domains that serve as a foundation for researchers approaching similar problems from different perspectives and disciplines (Bertrand and Wortmann 1981, Bertrand et al. 1990). A key assumption of both bodies of work – the stochastic inventory models for general supply chain networks and the linear programming model suggested in the handbook chapter and later explored in Spitter et al. (2005b) and Spitter et al. (2005a)- is the assumption of fixed, exogenous lead times that are independent of production quantities. This approach is widely used in both academic research and industrial practice, and has the advantage of yielding models that are both computationally more tractable and easier to explain than some alternatives. However, a case can be made that planning models treating production lead times as endogenous are worthy of study, although in a relatively early, often an exploratory, state of development. In this chapter I hope to convince the reader that workload-dependent lead times are an interesting research topic that yields interesting insights, and that while it is not obvious that they are always worth the effort needed to implement them, they may be useful in some contexts. I am deliberately adopting an informal, conversational tone in which to present the discussion, while providing sufficient references for the reader to be able to pursue any directions they may find of interest.

I shall assume a production system consisting of a general network of production units, whose internal operations are the responsibility of their local management and thus outside the control of, and invisible to, the planning function. The purpose of the planning function is to coordinate the flow of material across the production units to the end customer. This view can apply equally well to a supply chain, where the production units may represent different firms, or a single large complex facility, where the production units may represent specific workcenters or departments. I shall assume that the behavior of each production unit can be modelled to an acceptable degree of accuracy (however defined) as a queueing system, whose specific characteristics may vary with the environment under study. I will also assume the conventional discrete-time environment for production planning problems, where the planning horizon is divided into discrete periods of equal length, with decision and state variables being defined for each period. This is essentially the environment of the de Kok – Fransoo chapter, and I now discuss why lead times can be endogenous to a planning model.

3 Why Workload-Dependent Lead Times?

The purpose of the planning exercise is to ensure that production orders are released to the production units over time in a manner that ensures their completion as saleable products meeting customer demand in the "best" possible manner – minimum cost, maximum profit, or some other objective. In this context the cycle time (also referred to as flow time) of a production order, the time elapsing between its release to a production unit and the completion of its processing at that unit, moves to center stage. It is widely recognized that cycle times can be quite substantial, and certainly far from negligible, so for a given production order to be completed at a desired point in time, it must be released to production in advance of the demand it is to meet. This need is often expressed in terms of anticipation (Schneeweiss 2003) – the planning function, in order to make good decisions, must be able to anticipate, to some degree of accuracy, the impact of its release decisions on the shop-floor performance (specifically, the completion times of production orders) of the production units it is planning for.

Queueing theory (Buzacott and Shanthikumar 1993) has shown that the cycle time of a production order released at a specific point in time through a production unit is a random variable whose probability distribution is affected, among several other factors, by the utilization level of the resources in the production unit. However, the utilization level of resources is determined by the amount of work released to the production unit, which is the task of the planning model as we have defined it. Thus the outcome of the planning process at least partially determines the probability distribution of the cycle times of production orders released in a given period. This circularity has lain at the center of research into production-inventory systems since the inception of the field in the 1960s, and remains a significant challenge today.

I have discoursed upon cycle times, and have not yet mentioned lead time. Since decisions and system state are computed for each planning period, I will define the lead time associated with a production order released in a period as a parameter of the planning process that seeks to account for the cycle time the order will incur if released in that period. This lead time ought to be related to the probability distribution of the cycle time of the order released in that period, perhaps a percentile. It is likely to be prohibitively difficult to estimate probability distributions for the cycle time of each individual order released in a period which will depend on the other orders released in the same period. Thus, a simpler approach is to consider the population of all orders released to the production unit in a given period and consider a probability distribution of cycle times for a randomly selected order from this set. This results in a lead time associated with each production unit in each period as discussed in de Kok and Fransoo (De Kok and Fransoo 2003), although in their case the dependency on periods does not need to be considered since they assume time-stationary demand distributions. I shall refer to this type of lead time as planned lead times, noting they are exogenous to the planning model. The estimation of planned lead times for use in planning a specific production system is an interesting topic in itself that merits further study (Milne et al. 2015); in particular, the estimation of time-varying planned lead times involves a number of complications such as discontinuous jumps in lead times for consecutive periods that are discussed at length in Chapter 6 of Missbauer and Uzsoy (2020).

Exogenous planned lead times are by far the most common approach to anticipating the impact of cycle times in planning models, notably the widely used Material Requirements Planning (MRP) procedure and its variants (Baker 2003, Vollmann et al. 2005) and most linear and mixed integer production planning models (Hackman and Leachman 1989, Hackman 2008, Pochet and Wolsey 2006, Missbauer and Uzsoy 2020). While most commonly expressed as integer multiples of the planning period, the concept can easily be extended to fractional numbers of periods (Hackman and Leachman 1989, Missbauer and Uzsoy 2020). Extensive computational studies (Kacar et al. 2016) have shown that the use of fractional lead times can significantly enhance the performance of planning models over those using integer lead times only. Albey and Uzsoy (2015) have used simulation optimization to determine the planned lead times that maximize realized profit for a production system, and find that the use of fractional lead times can achieve performance comparable with that of much more complex clearing function models, supporting the findings of Kacar et al. (2016)).

It is helpful, in my opinion, to distinguish between two different uses of planned lead times. The first of these is the anticipation function in the sense of Schneeweiss (2003) – providing the planning function with an estimate of the consequences of its decisions on the performance of the production units it seeks to coordinate. The second is somewhat more subtle, in that planned lead times can serve as a mechanism for coordination between different agents involved in the production system. In the latter case, the planned lead times serve as performance targets that all concerned agree will be met with high probability, in a manner analogous to the treatment of replenishment lead times in the Guaranteed Service literature on inventory models (Eruguz et al. 2016). Thus the planned lead time allows any agent involved in the planning exercise to anticipate the behavior of production units with high confidence while planning its own operations, which is not the same thing as predicting the completion time of a production order. For instance, a production order completed early might well satisfy the planned lead time.

Given the above discussion I believe we need to consider two separate situations. If a plan is being computed for a number of production units with a centralized objective, and the results of this centralized plan can be communicated to all concerned, I see no reason why a planning model with workload-dependent, i.e., endogenous, lead times cannot be used since the planned completion times computed in the plan provide guidance for the production units. In the case of a decentralized environment, where a plan must be arrived at collaboratively by a number of autonomous actors each with their own objective and private information, the situation is less clear, and the value of exogenous planned lead times as a coordinating mechanism, essentially as a contract between parties specifying service guarantees, appears a lot more attractive.

Having attempted to make a case for the study, if not the immediate deployment, of planning models with workload-dependent lead times, I now provide a brief, informal discussion of the state of the art of these models and some of the important issues that need to be addressed for these to become more useful.

4 Modelling Workload-Dependent Lead Times

There have been three principal streams of work related to the development of planning models with workload-dependent lead times: iterative multimodel approaches, clearing functions and state-based models. I will not attempt a comprehensive coverage of the topic here, but rather focus on providing a brief conceptual overview, addressing some selected topics of interest.

4.1 Iterative Multimodel Methods

These approaches, whose origins date back to the work of Zäpfel (1984), decompose the planning problem into two subproblems: one that computes the optimal releases for a given set of planned lead times, and a second that estimates the realized cycle times those releases will lead to. The release planning problem is usually formulated as a linear program, while the cycle time estimation component has been addressed using simulation or, in a few cases, queueing models. Starting with an initial estimate of the planned lead times, at each iteration first the release planning subproblem and then the cycle time estimation subproblem are solved. The realized cycle times (or a related quantity) are then communicated to the release planning subproblem, and the next iteration begins. The iterations continue until some convergence criterion is satisfied. A variety of such models have been suggested over the years Missbauer and Uzsoy (2020), the most notable being those of Hung and Leachman (1996) and Kim and Kim (2001).

These approaches have the benefit of combining two well-known modelling approaches in an intuitive manner, but they face a number of difficulties. Computational experiments have shown that their convergence behavior can vary quite substantially across different methods (Irdem et al. 2010), the solution they converge to may depend on initialization, and they can be outperformed by clearing function models (Kacar et al. 2012). Most importantly, however, Missbauer (2020) has shown that procedures of this type that iterate on lead times behave as a price-based coordination mechanism that does not meet the theoretical requirements for convergence. Another issue with these approaches is the potentially high computational burden for large production systems due to the need for multiple simulation replications at each iteration. This last item may not be such a major issue in the future as it has been in the past, if the time-consuming simulation model can be replaced by a surrogate model such as a neural network trained offline, but the conceptual issues with convergence remain. The amount of data and computing power needed to train a reliable surrogate model for large production systems or supply chains is also a concern.

In summary, while this class of methods has the advantage of using well-known models in an intuitive way, making it relatively easy to implement them with today's software tools, their performance is not well understood and a lot more research is needed to bring them to a point where they can be used in practice, although such attempts have been made (Allison et al. 1997).

4.2 Clearing Functions

In its most common form (Missbauer and Uzsoy 2020), a clearing function describes a functional relation between the workload available to the resource in a planning period and its expected output during the period. Initially proposed in the late 1980s (Graves 1986, 1988, Srinivasan et al. 1988, Karmarkar 1989), this approach initially encountered difficulties in representing the behavior of multi-item systems, which were largely, although not completely, resolved by the Allocated Clearing Function formulation (Asmundsson et al. 2009). Similar flow-density functions representing the rate of traffic flow through a road segment as a function of the number of vehicles using the segment have been studied extensively in the traffic literature (Carey and Bowers 2012).

The principal advantage of this approach is that it provides a way of capturing the nonlinear relation between average workload and average throughput in an optimization model that, in the absence of lot-sizing considerations, can be solved efficiently using convex optimization solvers. In contrast to classical linear programming models with exogenous planned lead times, it also permits the calculation of informative dual prices for resources that are not fully utilized (Kefeli and Uzsoy 2016, Sutterer et al. 2022). Extensive computational experiments (Kacar et al. 2016, Haeussler et al. 2020) have shown that, especially when workloads vary over time, clearing function models can outperform LP models using fixed lead times, although allowing the use of fractional lead times and optimizing the exogenous lead times by simulation optimization (Albey and Uzsoy 2015) often yield comparable results. Ziarnetzky et al. (2015) find that the performance advantages of the clearing function approach are maintained in a rolling horizon environment. To the best of our knowledge this approach has yet to be implemented in practice. Detailed derivations and justification of this formulation are given in Asmundsson et al. (2009) and Chapter 7 of Missbauer and Uzsoy (2020).

While the clearing function approach has several attractive features, notably the ability to capture queueing behavior to at least a first degree of approximation in a tractable optimization model, a number of conceptual and practical difficulties remain to be addressed. The most common way of applying clearing functions to a production unit is to model each resource group (production unit) using its own clearing function that is independent of those at other stages in a manner equivalent to the widely used, but also widely questioned, decomposition assumption in queueing networks (Bitran and Tirupati 1988). A theoretically consistent clearing function can be derived for some steady-state queueing models, but experimental evidence suggests that these do not yield good performance when used without modification.

The problem of fitting clearing functions to experimental data is also fraught with difficulties (Gopalswamy and Uzsoy 2019). The explanatory variable of a clearing function obtained by fitting it to empirical or simulation data is the known, and hence deterministic, workload observed in the data, whereas its use in planning implies that the workload is a planned, and hence uncertain, quantity appropriately modeled as a random variable. This leads to biased estimates of output for a given workload per Jensen's Inequality which has proven difficult to resolve (Gopalswamy and Uzsoy 2019, Missbauer 2011). Other difficulties arise in the context of collecting empirical data for the purpose of fitting clearing functions. The basic approach is to simulate the production resource over time in the face of a particular set of releases, observing the workload and output in each planning period. However, the probability distribution of the system states, defined by the workload of each item in each period, depends on the specific releases used, directly affecting the estimate of the clearing function. This phenomenon is referred to as dataset shift in the data analytics vocabulary (Quinonero-Candela et al. 2008), and requires careful formulation of the process of clearing function estimation. An interesting and consistent observation (Kacar and Uzsoy 2015, Gopalswamy and Uzsoy 2019) is that rather than fitting the clearing function to optimize fit to the available data, fitting it to obtain the best performance from the production system using its decisions yields far better results. A fully satisfactory formulation of this problem remains to be developed. Finally, industrial data that could be used for fitting clearing functions is difficult to obtain, since most industrial facilities operate over a relatively limited range of product mixes and utilizations, rendering a validated simulation model essential for accessing a broad range of operating conditions. A number of authors (Fine and Graves 1989, Haeussler and Missbauer 2014) have found significant, qualitative differences between industrial and simulated data which are yet well understood. Clearly, building a validated simulation model of a large facility or supply chain is a significant investment of resources in its own right. How to combine limited industrial data with simulation data is also an interesting research question.

4.3 State-Based Models

Initially proposed by Omar et al. (2017), this class of models constitutes the most recent development in the attempt to capture the behavior of production-inventory systems. The basic approach is to assume the system evolves through a (potentially very large) set of discrete states, each of which is associated with an expected level of output. The state is usually defined by the WIP of each product in the system, measured at discrete time points t which can be viewed as defining discrete planning periods. Given the system state at the start of some period t, the planning problem seeks to determine release and production actions that will bring the system into a desired state at the start of period $t+1$. This class of models have generally been referred to as "data driven" models. However, given the broad association of the term "data driven" with statistical learning and deep learning models, I prefer to refer to them as "state-based" models to avoid possible confusion.

The basic idea of these models is to represent the behavior of a production system as evolving through a finite, discrete set of states defined by discrete time periods as in most of the other models above. The states are usually defined based on some measure of the available WIP of each product, and each state is associated with an expected output of each product at the end of the current planning period. The production planning problem can then be formulated as that of selecting exactly one state in each period for the system, with state transitions governed by the material balance equations for WIP and finished inventory as well as other relevant constraints, most notably capacity constraints. Release decisions and production quantities can then be inferred from the beginning and ending states for the planning period.

Since the system states are defined for the entire production system without considering individual production units or resources, these models tend to have fewer decision variables and constraints that the others discussed above. The obvious difficulty, which these models share with a number of stochastic models such as Markov Decision Processes, is the very large number of states required to describe the behavior of a large production system. Thus implementing these models requires identifying a suitably representative subset of the states that will, on the one hand, permit the computation of sufficiently accurate production plans while allowing the optimization model, which is a mixed-integer program, to be computationally tractable. Like the clearing function models discussed above, it also requires a large amount of off-line computation to determine the expected output for each of the states considered. These computations can be performed using queueing models or simulation. Another potential difficulty, addressed by Völker and Mönch (2023), is that large production systems like semiconductor wafer fabs may have cycle times considerably longer than one planning period.

Research on these models is in its very early stages, but computational experiments (Gopalswamy and Uzsoy 2018, Omar et al. 2017, Völker and Mönch 2023) have found that when appropriately parameterized state-driven models can yield solutions comparable in quality to those obtained by the clearing function models.

5 Problems with Stochastic Demand

It has always been a source of curiosity to me that stochastic inventory theory and production planning have historically developed largely independently of each other. After all, inventory has to be produced at some point! Inventory theory has shown that the amount of inventory necessary to maintain a specified level of customer service is driven by the probability distribution of the demand over the replenishment lead time, the time elapsing between an order replenishing the inventory in question being placed and which, if inventory is replenished from a production system, corresponds to the time between the order being communicated to the production unit and its becoming available to meet demand – precisely the cycle time we have been discussing above. The same arguments relating the probability distribution of the cycle time to resource utilization, determined in this case by the size of the order placed relative to the capacity of the replenishing production unit are now valid.

However, in the case of stochastic demand, there is more. Classical results from inventory theory (Axsater 2010, Zipkin 2000) have shown that that the mean and variance of the lead time demand are driven by the mean and variance of the demand per period and the lead time; queueing tells us that the mean and variance of the cycle time increase with resource utilization. So we now have a version of the circularity related to cycle time for safety stocks: producing additional material for safety stocks will increase the mean and variance of the lead time demand, requiring even more safety stock. This suggests that in an ideal situation, safety stocks should also be endogenous to planning models along with lead times.

Despite the extensive development of stochastic optimization methods over the last two decades, these tools do not appear to have made deep inroads into the domains of production and inventory planning. The principal approaches one would consider are multistage stochastic programming (Birge and Louveaux 1997) and robust optimization (Bertsimas and Sim 2004, Bertsimas and Thiele 2006). Clearing functions to represent workload-dependent lead times can be incorporated into these models quite directly, and computational studies for simple, single-stage single-item systems (Aouam and Uzsoy 2015) have proven promising, noting that the way in which uncertainty is modelled in the different approaches requires quite different approaches to setting the parameters for these models.

Another body of work (Albey et al. 2015, Ziarnetzky et al. 2018, 2020) has sought to develop approximate solutions to the problem with stochastic demand by formulating mathematical programming models combining clearing functions with chance constraints seeking to capture desired service levels. Stochastic demand is modelled using the Martingale Model of Forecast Evolution (MMFE) (Heath and Jackson 1994, Graves et al. 1998) which represents the evolution of demand forecasts over time as additional information is obtained. Ziarnetzky et al. (2020) conduct extensive simulation experiments for a large semiconductor manufacturing facility in a rolling horizon environment and find that the models with workload-dependent lead times generally outperform those with fixed lead times.

In summary, a variety of methods involving workload-dependent lead times have been examined for planning problems with stochastic demand, but this work remains at an early stage of development, and offers many opportunities for interesting research that would have relevance to other areas of supply chain management.

6 Summary and Conclusions

In this chapter I have tried to provide a high-level overview of planning models with workload-dependent lead times. While happily acknowledging the long history and several advantages of models with exogenous planned lead times, there are interesting and difficult problems remaining even for this well-trodden area, such as how to estimate the planned lead times that lead to the best performance of the planning model in which they are used, as opposed to the fitting the empirical data best. Each of the three different approaches currently available to capture workload-dependent lead times has its pros and cons, but all of them are in essentially an exploratory stage of development and requiring substantial additional research to understand their behavior and address some of their known weaknesses. An important question is under what conditions the use of workload-dependent lead times in planning models yields sufficient benefits over more conventional approaches to justify their additional complexity and computational requirements.

From Equations to Solutions: A Journey of Academia-Industry Partnerships

Tom Van Woensel

Abstract: This chapter delves into Ton's specific approach to research, emphasizing the importance of understanding real-world challenges to develop impactful theories. Through his work at Eindhoven University of Technology (TU/e) and as Director of CWI Amsterdam, Ton has demonstrated that "There is nothing as practical as a good theory." when solving complex industrial problems. The chapter highlights his contributions to supply chain management, collaborative partnerships, and the European Supply Chain Forum (ESCF), showcasing how academic-industrial collaborations can drive innovation and create tangible benefits for both sectors.

1 Introduction

In the deterministic world of academic research, where variables behave themselves and uncertainty is often swept under the rug, two scholars dared to venture into the chaotic realm of reality. It was 2003 when I first crossed paths with the inimitable Ton de Kok, a man whose career had been built on taming the wild beast of stochasticity in supply chain management.

Picture a world where delivery trucks always arrive on time, customers know exactly what they want (and when they want it), and inventory levels are as predictable as a metronome. Sounds delightful. Well, it's about as accurate as a unicorn riding a hoverboard. Ton and I recognized that the messy, unpredictable nature of the real world was being underrepresented in transportation and logistics research.

Thus began our joint academic journey, the supply chain stochasticity guru and the transportation optimization enthusiast, united in our quest to inject a healthy dose of reality into logistics research. Our mission? To boldly go where few researchers had gone before: into the heart of uncertainty itself. We set out to tackle the wild world of time-dependent and stochastic travel times, uncertain demands, and all the other unpredictable elements that make logistics professionals wake up in a cold sweat.

So, prepare yourself for a journey through the fascinating career of Ton de Kok - a man who has spent decades proving that in the world of supply chain management, the only certainty is uncertainty itself. From his early days at Philips Electronics to his reign as the stochastic sovereign of the OPAC group, Ton's career is a testament to the power of embracing chaos and turning it into a competitive advantage.

Buckle up because, in the world of Ton de Kok, the ride is always bumpy, and the destination is often uncertain, but the insights are guaranteed to be groundbreaking.

2 The Academic-Industrial Interface: A Career of Bridging Theory and Practice

Ton de Kok's career is a testament to the power of synergy between academic research and industrial practice in operations management. Like a skilled tightrope walker, de Kok has masterfully balanced on the wire between the ivory tower of academia and the concrete jungle of industry, occasionally wobbling but never falling.

His journey began at Philips Electronics, where he honed his consultancy and process innovation skills, laying the foundation for his unique research approach. Little did Philips know that they were nurturing a future supply chain superhero with the power to optimize multi-echelon inventory systems faster than a speeding bullet.

In 1991, de Kok moved to academia, initially as a part-time professor in Industrial Mathematics at Eindhoven University of Technology (TU/e). By 1992, he had gone full-time, beginning a prolific academic career. His industrial background proved invaluable, allowing him to bridge the gap between theoretical models and real-world challenges. It was as if he had one foot in the world of abstract mathematics and the other in a warehouse full of inventory, performing a complex dance of optimization.

Throughout his tenure at TU/e, de Kok maintained strong ties with industry, refusing to be confined to the ivory tower. From 1996 to 2018, he served as the Founding Director of the European Supply Chain Forum (ESCF), fostering collaboration between academia and leading companies. This platform enabled him to stay connected with current industry challenges and ensure the relevance of his research. It was like having a permanent stethoscope on the beating heart of the industry, allowing him to diagnose and treat supply chain ailments with academic precision.

De Kok's commitment to practical applications led him to undertake a sabbatical at ASML, the world leader in semiconductor manufacturing equipment. During this period, he developed innovative solutions for supply chain planning, further demonstrating the value of academic-industrial partnerships. It was as if he had taken a deep dive into the semiconductor industry and emerged with a silicon chip of supply chain wisdom.

In 2018, de Kok co-founded ChainScope, a start-up company that leverages his research on multi-item, multi-echelon inventory systems. As the Chief Technology Officer, he continues to apply his academic insights to solve complex supply chain problems for businesses. ChainScope, developed by de Kok, is a software tool designed to optimize Supply Chain Operations Planning (SCOP), aiming to coordinate the release of materials and resources in the supply network to meet customer service constraints at minimal cost. It's like having a supply chain crystal ball, but one backed by rigorous mathematical models.

The culmination of de Kok's career came in 2020 when he was appointed Director of the Centrum Wiskunde & Informatica (CWI) in Amsterdam, the Netherlands' national mathematics and computer science research institute. This role allows him to further promote academic research and industrial innovation integration nationally, like a matchmaker for math and industry.

Throughout his career, de Kok has consistently emphasized the importance of translating theoretical concepts into practical solutions. His work in quantitative analysis has significantly impacted various aspects of operations management, particularly in inventory control and supply chain optimization. De Kok's research has led to the development of algorithms for computing inventory control parameters, which have been implemented in over 20 companies. These implementations, often resulting from MSc projects, demonstrate the practical applicability of his theoretical work.

One of de Kok's most notable contributions is in multi-echelon inventory systems. He developed new evaluation and optimization methods that have significantly advanced the understanding and management of complex supply chains. These methods have proven particularly valuable in Configure-To-Order supply chains and the process industry, which are his research focus areas.

However, de Kok's journey was not without challenges. Bridging the gap between academia and industry often meant speaking two languages: the precise, jargon-filled mathematics dialect and the profit-driven business vernacular. He had to become fluent in both, often as a translator between these two worlds. Moreover, the pace of academic research and the urgency of industrial needs don't always align, requiring careful balancing and prioritization.

De Kok managed these challenges by maintaining a foot in both worlds. His roles at ESCF and ChainScope and his sabbaticals allowed him to stay attuned to industry needs, while his academic position provided the rigor and resources for in-depth research. He embraced the tension between theory and practice, using it as a driving force for innovation rather than a source of conflict.

In his words, "Over the years, I have found that our research helps translate today's exceptions into tomorrow's routines. That leaves more time for people in practice to reflect and work on challenging projects". This approach has not only advanced the field of operations management but has also had a tangible impact on industry practices, exemplifying the power of bridging academia and industry.

Through his work with ChainScope and other quantitative tools, de Kok has consistently shown that rigorous mathematical analysis can significantly improve real-world supply chain operations. His approach combines the strength of mathematically determined inventory control policies with the tacit knowledge of human decision-makers, providing a robust framework for addressing complex operational challenges. It is as if he created a supply chain Swiss Army knife equipped with tools for every logistical challenge.

Ton de Kok's career is a masterclass in academic-industrial symbiosis. He has shown that with the right approach, the theoretical rigor of academia and the practical demands of industry can create a powerful synergy, driving innovation and solving real-world problems. His journey from a mathematician to a supply chain guru inspires future generations of researchers and practitioners alike, proving that sometimes, the best way to solve a problem is to have a foot in both worlds - even if it means doing a constant mathematical split.

3 The European Supply Chain Forum: A Cornerstone of Industry-Academia Collaboration

The European Supply Chain Forum (ESCF), established in 1996 by Professor Ton de Kok, is a groundbreaking example of structured industry-academia collaboration in supply chain management. As the founding director, de Kok established a robust platform to connect theoretical research with practical application.

The ESCF's origins can be traced to a two-day meeting organized by TU/e academics, bringing together 30 leading global players to discuss supply chain management challenges. From this initial gathering, de Kok conceptualized a forum where academia and industry could jointly address practical and scientific supply chain issues through technology, talent development, and innovative thinking.

Under de Kok's leadership from 1996 to 2018, ESCF flourished into a dynamic ecosystem. The forum's structure solidified in 1998, featuring four annual workshops on cutting-edge supply chain topics and an annual forum meeting. An advisory board, including executives from multinational companies, was formed to guide the forum's direction.

De Kok's vision for ESCF eventually led to a three-pillar model, formalized in 2020: Talent, Knowledge, and Network. The Talent pillar connected companies with top students and supply chain professionals from TU/e. The Knowledge pillar focused on generating and exchanging scientific and practical supply chain knowledge. The Network pillar fostered a trusted environment for members to share best practices and engage in peer-to-peer contacts.

The impact of ESCF under de Kok's leadership was substantial. From 2001 to 2011, research conducted by ESCF in collaboration with its members yielded 1 billion euros in benefits. This impressive figure underscores the forum's success in translating academic insights into real-world solutions. Operating with an annual budget of approximately 1 million euros during this period, the forum generated a thousand-fold return for its members. The billion-euro yield was not merely theoretical but represented tangible benefits for member companies, including cost savings, increased efficiency, and new revenue streams from implementing innovative supply chain strategies and technologies (De Kok 2013).

Throughout de Kok's tenure, ESCF consistently explored emerging topics with its members, leading to inspirational events and shaping the European supply chain landscape by promoting sustainable and circular economy practices. The success of ESCF today is mainly attributable to the strong foundations laid by de Kok. His vision of creating value in supply chains through collaborative efforts continues to guide the forum, serving as a model for effective academic-industrial partnerships and demonstrating how academic rigor and industrial practicality can drive significant advancements in supply chain management practices.

4 Overcoming Challenges in Academic-Industrial Partnerships

The 4D model, Design, Deliver, Disseminate, and Demonstrate, offers a structured approach for bridging the gap between academic research and industry practice in operations management, supply chain management, and logistics. This model addresses the significant challenge of aligning academic research with industry needs, as highlighted by the divergent priorities and perspectives of academic researchers and industry practitioners (Dwivedi et al. 2024a,b).

Design: This phase involves carefully planning research projects with industry relevance in mind. Researchers can collaborate with industry partners to identify pressing challenges, such as optimizing production schedules, improving inventory management, or enhancing supply chain resilience. The Design phase emphasizes the importance of understanding industry needs and agreeing on how impact will be interpreted and measured by different stakeholders (Dwivedi et al. 2024a). It includes identifying key business problems and framing research questions to address them.

Deliver: The Deliver stage focuses on conducting rigorous research and producing high-quality outputs that can be applied in real-world contexts. This might involve developing new algorithms for vehicle routing problems that consider cost efficiency and environmental impact in logistics. This stage aims to convey research findings in a form and structure that can be easily understood by all parties (Dwivedi et al. 2024a).

Disseminate: Dissemination is crucial for ensuring that research findings reach the intended audience in the industry. Researchers must translate their findings into language and formats accessible to industry practitioners. This could involve creating user-friendly software tools or decision support systems based on academic research. The Disseminate phase may involve translating academic jargon into more accessible language and utilizing various communication mechanisms such as open access and multi-media platforms (Dwivedi et al. 2024a).

Demonstrate: This phase showcases the research's tangible benefits and practical applications, often through case studies or pilot implementations. In operations management, this could mean implementing new inventory control systems in partner companies and measuring efficiency and cost reduction improvements.

This phase also serves as an evaluation function to ensure that the research findings and recommendations are realized and measured in terms of their impact and changes to policy (Dwivedi et al. 2024a).

The 4D model aims to address several key challenges in academic-industrial partnerships:

1. It helps overcome the misalignment of priorities between academia and industry.

2. It ensures research outputs are more accessible and relevant to practitioners.

3. It provides a framework for measuring and demonstrating impact beyond traditional academic metrics.

By following this 4D model, researchers in operations management, supply chain management, and logistics can ensure that their work advances theoretical knowledge and directly impacts industry practice. This approach can lead to more robust solutions for complex challenges in these fields, such as managing multi-echelon inventory systems, optimizing last-mile delivery in urban environments, or developing more sustainable supply chain practices.

Professor Ton de Kok's career exemplifies the successful implementation of this approach in bridging academia and industry. His work at the European Supply Chain Forum (ESCF) serves as a "permanent thermometer" for industry needs, allowing him to design research projects that address real-world challenges. De Kok's sabbatical at ASML demonstrates the Deliver phase, where he applied his academic expertise to develop innovative solutions for supply chain planning in a high-tech manufacturing environment.

In the Disseminate phase, de Kok was particularly effective through his leadership of the ESCF, creating a platform for knowledge exchange between academia and industry. His ability to translate complex theoretical concepts into practical applications, as seen in his work with ChainScope, showcases the Demonstrate phase of the 4D model.

De Kok's career-long commitment to bridging theory and practice aligns closely with the principles of the 4D model. His success in fostering academic-industrial partnerships and delivering impactful research demonstrates the potential of this approach in overcoming the traditional barriers between academia and industry in operations management. The 4D model should be considered as a "co-productive" approach whenever possible, with both academics and practitioners participating in activities underlying the model, enabling a clear identification of the eventual impact of the proposed research and an appropriate design to achieve such an impact (Dwivedi et al. 2024a).

5 Our Collaborative Research

Our collaborative research has significantly advanced sustainable logistics and urban distribution. It addresses complex challenges in modern supply chains, particularly in urban environments where efficiency and sustainability are paramount. Our partnership has led to innovative approaches to vehicle routing problems (VRPs), incorporating environmental considerations, stochasticity, and uncertainty into traditional routing optimization models.

One key contribution of our joint work is the development of sophisticated models that integrate time-dependency and uncertainty into Vehicle Routing Problems (VRPs). Our research has focused on addressing the complex challenges faced by real-world logistics operations, where fluctuating travel times, variable customer demands, and unpredictable service times significantly impact efficiency and reliability.

Recognizing that traditional static VRP models often fall short in dynamic urban environments, we developed stochastic VRP models that can handle these uncertainties. Our work has resulted in more robust and realistic solutions for industry practitioners, enhancing the reliability and efficiency of logistics operations. For instance, we have developed models for time-dependent travel times, allowing for more accurate route planning in congested urban areas where travel speeds vary significantly throughout the day. Furthermore, our research has explored the concept of reliability in VRPs, addressing the critical need for consistent and dependable delivery schedules in the face of uncertainties. By incorporating these real-world constraints and uncertainties, we also significantly advanced the practical applicability of VRP solutions in complex, dynamic logistics environments.

Our collaboration has also made substantial contributions to the emerging field of city logistics, proposing innovative two-echelon distribution systems that combine urban consolidation centers with environmentally friendly last-mile delivery vehicles. This approach has proven particularly effective in addressing the unique challenges of urban distribution, where congestion, environmental concerns, and customer expectations intersect.

Collaborations with talented PhD students have enriched the research, bringing fresh perspectives and contributing significantly to the research output.

As detailed in her thesis "Time and Timing in Vehicle Routing Problems," Ola Jabali's work explored the critical aspects of time-dependent travel times and self-imposed time windows in VRPs. Jabali's research introduced the time-dependent VRP, which is designed to handle unexpected delays at customer locations. It developed solution methods combining disruptions in a Tabu Search procedure (Jabali 2010).

Said Dabia's thesis, "Time and Multiple Objectives in Scheduling and Routing Problems," expanded on time dependency in VRPs. Dabia's work focused on developing multi-objective optimization models that balance travel time, cost, and environmental impact in routing decisions, further enhancing the practical applicability of VRP solutions in complex real-world scenarios (Dabia 2012).

Duygu Tas's research, outlined in the thesis "Time and Reliability in Vehicle Routing Problems," addressed the critical issue of reliability in logistics operations. Tas's work explored methods to improve the consistency and dependability of delivery schedules in the face of uncertainties, contributing to developing more robust routing strategies (Tas 2013).

Derya Sever's thesis, "Routing in Stochastic Networks," delved into the challenges of navigating uncertain and dynamic transportation networks. Sever's research focused on developing algorithms and models capable of adapting to changing network conditions, such as traffic congestion or unexpected road closures, further enhancing the resilience of routing solutions in real-world applications (Sever 2014).

Our collaborative efforts resulted in substantial work, including publications in top-tier journals and presentations at international conferences. Our research has advanced the academic understanding of VRPs and provided practical solutions for industry practitioners, helping shape the future of sustainable and efficient logistics operations in urban environments.

6 Some Concluding Words...

As we conclude this journey through Ton de Kok's remarkable career and our collaborative efforts, I reflect on our partnership's profound impact. I truly appreciated the collaboration with Ton, which has been nothing short of transformative. Like a master chef combining unexpected ingredients, Ton's ability to blend rigorous academic theory with practical industry insights has created a recipe for success that continues to inspire.

Over the years, I learned an immense amount from Ton. His approach to problem-solving is like watching a chess grandmaster at work, always thinking several moves ahead and considering the implications of each decision. From him, I gained not just knowledge but a whole new perspective on how to approach complex supply chain challenges. The legacy of our work extends far beyond our contributions. The strong industry-academia collaboration we fostered will continue to thrive within the OPAC group and certainly in my research.

Of course, the challenges in this collaboration remain, as highlighted by the 4D model. Bridging the gap between academic rigor and industrial practicality is often like trying to build a bridge while simultaneously standing on both sides of the river. It requires constant effort, communication, and a willingness to adapt. But it is precisely these challenges that make the work so rewarding.

This approach, I believe, is what good Industrial Engineering research should be all about. It is not just about solving equations or optimizing processes but about making a tangible difference in the real world. It is about turning theoretical models into practical solutions that can improve efficiency, reduce costs, and ultimately make businesses successful.

As we look to the future, I am excited about the possibilities. The seeds planted will continue to grow, and I'm confident that the next generation of researchers and practitioners will take this work further. They will face new challenges, but armed with the tools and mindset we've developed, they will be well-equipped to tackle whatever comes their way. In the spirit of Ton's love for practical applications, I conclude with a little experiment: the first person who reads this text thoroughly and reaches out to me will get a bottle of wine. After all, what better way to celebrate the fruits of our research than with a good vintage?

I look forward to the future of supply chain management on the interface between academia and practice. May it be as rich, complex, and satisfying as the collaborations that drive it forward.

Nawoord

Als samenstellers van dit liber amicorum hebben wij met veel plezier en bewondering de uiteenlopende bijdragen gelezen. Wij hebben ervoor gekozen om de hoofdstukken te rangschikken in de volgorde waarin de auteurs Ton in de loop der jaren hebben leren kennen: van degenen die hem de eerste richting gaven in het vakgebied, tot degenen die hij op zijn beurt voor het vak interesseerde en die hij wist te enthousiasmeren met zijn manier van werken. Zo ontvouwt zich, min of meer chronologisch, een beeld van Ton als mens, én als onderzoeker.

De ontwikkeling van de synchronized base stock policies door Ton laat zien hoe er, door schakels op de juiste manier in beweging te brengen, een systeem ontstaat dat als geheel presteert zoals geen enkele schakel afzonderlijk zou kunnen. Tons eigen rol in het onderzoeksveld is in zekere zin een illustratie van dit principe: Zijn vermogen om vele mensen voor het belang van relevant supply chain onderzoek te enthousiasmeren, bracht een keten tot stand waarin brede samenwerking en inspiratie elkaar versterken.

Zelf bevinden wij ons aan het eind van de keten: wij behoren tot de laatsten die Ton in zijn loopbaan heeft leren kennen. Ook wij hebben ervaren hoe zijn energie en betrokkenheid aanstekelijk werken. Met dit boek hopen we iets terug te geven van wat hij voor zo velen heeft betekend. We zijn de auteurs dankbaar voor hun inspirerende bijdragen, en bovenal Ton — voor zijn intellect, zijn warmte en zijn unieke vermogen om mensen en ideeën samen te brengen.

Zümbül Atan en Willem van Jaarsveld

Bibliography

I. Adan, M. van Eenige, and J. Resing. Fitting discrete distributions on the first two moments. *Probability in the Engineering and Informational Sciences*, 4:623–632, 1995.

E. A. Albey and R. Uzsoy. Lead time modeling in production planning. In *2015 winter simulation conference (wsc)*. IEEE, 2015.

E. A. Albey, A. Norouzi, K. G. Kempf, and R. Uzsoy. Demand modeling with forecast evolution: An application to production planning. *IEEE Transactions on Semiconductor Manufacturing*, 28(3):374–384, 2015.

P. Alfredsson and J. Verrijdt. Modeling emergency supply flexibility in a two-echelon inventory system. *Management Science*, 45:1416–1431, 1999.

R.A.H. Allison, J. Yu, L.H. Tsai, C. Liu, M. Drummond, D. Kayton, T. Sustae, and J. Witte. Macro model development as a bridge between factory level simulation and lp enterprise systems. In *Twenty First IEEE/CPMT International Electronics Manufacturing Technology Symposium Proceedings 1997 IEMT Symposium*, pages 408–416. IEEE, 1997.

A. Andriolo, D. Battini, R. W. Grubbström, A. Persona, and F. Sgarbossa. A century of evolution from Harris's basic lot size model: Survey and research agenda. *International Journal of Production Economics*, 155:16–38, 2014. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2014.01.013. URL `https://www.sciencedirect.com/science/article/pii/S0925527314000152`. Celebrating a century of the economic order quantity model.

T. Aouam and R. Uzsoy. Zero-order production planning models with stochastic demand and workload-dependent lead times. *International Journal of Production Research*, 53(6):1–19, 2015.

J. M. Asmundsson, R. L. Rardin, C. H. Turkseven, and R. Uzsoy. Production planning with resources subject to congestion. *Naval Research Logistics*, 56(2):142–157, 2009.

S. N. Atadeniz and S. V. Sridharan. Effectiveness of nervousness reduction policies when capacity is constrained. *International Journal of Production Research*, 58 (13):4121–4137, 2020.

Z. Atan, A. G. De Kok, N. P. Dellaert, R. van Boxel, and F. Janssen. Setting planned leadtimes in customer-order-driven assembly systems. *Manufacturing & Service Operations Management*, 18(1):122–140, 2016.

S. Axsäter. *Inventory Control, 2nd edition.* Springer, 2006.

S. Axsater. *Inventory Control.* Springer, New York, 2010.

K. R. Baker. *Requirements Planning*, volume 11 of *Handbooks in Operations Research and Management Science*, pages 597–675. Elsevier, Amsterdam, 2003.

J. W. M. Bertrand and J. C. Wortmann. *Production control and information systems for component-manufacturing shops.* Elsevier, Amsterdam, 1981.

J. W. M. Bertrand, J. C. Wortmann, and J. Wijngaard. *Production control: a structural and design oriented approach.* Elsevier Science Inc., 1990.

D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.

D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, 2006.

J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming.* Springer, New York, 1997.

G. R. Bitran and D. Tirupati. Multiproduct queueing networks with deterministic routing: Decomposition approach and the notion of interference. *Operations Research*, 34(1):75–100, 1988.

J. D. Blackburn, D. H. Kropp, and R. A. Millen. A comparison of strategies to dampen nervousness in mrp systems. *Management science*, 32(4):413–429, 1986.

J. A. Buzacott. The effect of station breakdowns and random processing times on the capacity of flow lines with in-process storage. *AIIE Transactions*, 4(4):308–312, 1972.

J. A. Buzacott and J. G. Shanthikumar. *Stochastic models of manufacturing systems.* Prentice-Hall, Englewood Cliffs, 1993.

W.E. Camp. Determining the production order quantity. *Management Engineering*, 2: 17–18, 1922.

M. Carey and M. Bowers. A review of properties of flow–density functions. *Transport Reviews*, 32(1):49–73, 2012.

S. Chopra and P. Meindl. *Supply Chain Management, 6th global ed.* Pearson, 2016.

A. J. Clark and H. Scarf. Optimal policies for a multi-echelon inventory problem. *Management science*, 6(4):475–490, 1960.

S. Dabia. *Time and multiple objectives in scheduling and routing problems.* Phd thesis, Industrial Engineering and Innovation Sciences, 2012.

J. N. Darroch. On the traffic-light queue. *Annals of Mathematical Statistics*, 35:380–388, 1964.

A. G. De Kok. A moment-itration method for approximating the waiting-time characteristics of the gi/i/1 queue. *Prob. Eng. and Inf. Sc.*, 3:273–287, 1989.

A. G. de Kok. A moment-iteration method for approximating the waiting-time characteristics of the $gi/g/1$ queue. *Probability in the Engineering and Informational Sciences*, 3:273–287, 1989.

A. G. De Kok. Computationally efficient approximations for balanced flowlines with finite intermediate buffers. *International Journal of Production Research*, 28(2): 401–419, 1990a.

A. G. De Kok. Hierarchical production planning for consumer goods. *European Journal of Operational Research*, 45(1):55–69, 1990b.

A. G. De Kok. Capacity allocation and outsourcing in a process industry. *International Journal of Production Economics*, 68:229–239, 2000.

A. G. De Kok. Cursor, 2013. URL `https://issuu.com/tuecursor/docs/cursor_6_nl`. Accessed on January 15, 2025.

A. G. De Kok. Buffering against uncertainty in high-tech supply chains. In *2015 winter simulation conference (wsc)*, pages 1–10. IEEE, 2015.

A. G. De Kok. Modelling short-term manufacturing flexibility by human intervention and its impact on performance. *International Journal of Production Research*, 56 (1-2):447–458, 2018a.

A. G. De Kok. Inventory management: Modeling real-life supply chains and empirical validity. *Foundations and Trends® in Technology, Information and Operations Management*, 11(4):343–437, 2018b.

A. G. De Kok and J. C. Fransoo. *Planning supply chain operations: Definition and comparison of planning concepts*, volume 11 of *Handbooks in operations research and management science*, pages 597–675. Elsevier, Amsterdam, 2003.

A. G. De Kok and S. C. Graves, editors. *Supply chain management: Design, coordination and operation*, volume 11. Elsevier, Amsterdam, 2003.

A. G. De Kok and K. Inderfurth. Nervousness in inventory management: comparison of basic control rules. *European Journal of Operational Research*, 103(1):55–82, 1997.

A. G. De Kok and Jeremy W. C. H. Visschers. Analysis of assembly systems with service level constraints. *International Journal of Production Economics*, 59(1-3):313–326, 1999.

A. G. De Kok, H. C. Tijms, and F. A. Van der Duyn Schouten. Approximations for the single-product production-inventory system with compound poisson demand and service-level constraints. *Advances in Applied Probability*, 16:378–401, 1984.

A. G. De Kok, F. Janssen, J. Van Doremalen, E. Van Wachem, M. Clerkx, and W. Peeters. Philips electronics synchronizes its supply chain to end the bullwhip effect. *Interfaces*, 35(1):37–48, 2005.

T. De Kok, J. Van Dalen, and J. Van Hillegersberg. *Cross-chain collaboration in the fast moving consumer goods supply chain.* Eindhoven University of Technology, 2015.

M. B. M. De Koster. Approximate analysis of production systems. *European journal of operational research*, 37(2):214–226, 1988.

R. Dekker, J.B.G. Frenk, M.J. Kleijn, and A.G. de Kok. On the newsboy problem with a cutoff transaction size. *IIE TRansactions*, 32(5):461–469, 2000. doi: 10.1023/A:1007649027857. URL `https://doi.org/10.1023/A:1007649027857`.

M. I. Dessouky, N. Moray, and B. Kijowski. Strategic behavior and scheduling theory. *Human Factors*, 37(3):443–472, 1995.

E. B. Diks and A. G. De Kok. Optimal control of a divergent multi-echelon inventory system. *European journal of operational research*, 111(1):75–97, 1998.

E. B. Diks and A. G. De Kok. Computational results for the control of a divergent n-echelon inventory system. *International Journal of Production Economics*, 59(1-3): 327–336, 1999.

E. B. Diks, A. G. De Kok, and A. G. Lagodimos. Multi-echelon systems: A service measure perspective. *European Journal of Operational Research*, 95(2):241–263, 1996.

M.K. Doğru. *Optimal Control of One-Warehouse Multi-Retailer Systems: An Assessment of the Balance Assumption*. Phd thesis, Industrial Engineering and Innovation Sciences, 2005.

Y. K. Dwivedi, L. Hughes, and A. Jeyaraj. Academic research and industry need more mutual feedback. LSE Impact of Social Sciences blog, March 2024a. Accessed on January 15, 2025.

Y. K. Dwivedi, A. Jeyaraj, L. Hughes, G. H. Davies, M. Ahuja, M. A. Albashrawi, A. S. Al-Busaidi, S. Al-Sharhan, K. I. Al-Sulaiti, and L. Altinay. Real impact: Challenges and opportunities in bridging the gap between research and practice – making a difference in industry, policy, and society. *International Journal of Information Management*, 78(102750), 2024b.

G. Eppen and L. Schrage. *Centralized Ordering Policies in a Multi-Warehouse System with Lead Times and Random Demand*, volume 16 of *Studies in the Management Sciences*, pages 51–67. North-Holland, Amsterdam, 1981.

A. S. Eruguz, E. Sahin, Z. Jemai, and Y. Dallery. A comprehensive survey of guaranteed-service models for multi-echelon inventory optimization. *International Journal of Production Economics*, 172:110–125, 2016.

C. H. Fine and S. C. Graves. A tactical planning model for manufacturing subcomponents of mainframe computers. *Journal of Manufacturing and Operations Management*, 2:4–34, 1989.

P. W. Frenken. Decomposition algorithms with ept based input: A case study in the automobile industry. Master's thesis, Eindhoven University of Technology, 2007.

J. R. Galbraith. Organization design: An information processing view. *Interfaces*, 4(3): 28–36, 1974.

S. B. Gershwin. An efficient decomposition method for the approximate evaluation of tandem queues with finite storage space and blocking. *Operations Research*, 35: 291–305, 1987.

E. M. Goldratt. Optimized production timetable: a revolutionary program for industry. In *APICS 23rd Annual Conference Proceedings*, pages 14–17, 1980.

K. Gopalswamy and R. Uzsoy. An exploratory comparison of clearing function and data-driven production planning models. In *2018 winter simulation conference (wsc)*. IEEE, 2018.

K. Gopalswamy and R. Uzsoy. A data-driven iterative refinement approach for estimating clearing functions from simulation models of production systems. *International Journal of Production Research*, 57(19):6013–6030, 2019.

J. Göttgens and R. Dekker. Operations research at a dutch pharmacy. *Medium Econometrische Toepassingen*, 15(4):8–14, 2007.

S. C. Graves. A tactical planning model for a job shop. *Operations Research*, 34(4): 522–533, 1986.

S. C. Graves. Safety stocks in manufacturing systems. *Journal of Manufacturing and Operations Management*, 1:67–101, 1988.

S. C. Graves. A multiechelon inventory model with fixed replenishment intervals. *Management Science*, 42(1):1–18, 1996.

S. C. Graves. How to think about planned lead times. *International Journal of Production Research*, 60(1):231–241, 2022.

S. C. Graves and S. P. Willems. Optimizing strategic safety stock placement in supply chains. *Manufacturing & Service Operations Management*, 2(1):68–83, 2000.

S. C. Graves, D. B. Kletter, and W. B. Hetzel. Dynamic model for requirements planning with application to supply chain optimization. *Operations Research*, 46(3):35–49, 1998.

S. T. Hackman. *Production Economics: Integrating the Microeconomic and Engineering Perspectives*. Springer, Berlin, 2008.

S. T. Hackman and R. C. Leachman. A general framework for modeling production. *Management Science*, 35(4):478–495, 1989.

G. Hadley. Analysis of inventory systems, 1963.

S. Haeussler and H. Missbauer. Empirical validation of meta-models of workcentres in order release planning. *International Journal of Production Economics*, 149:102–116, 2014.

S. Haeussler, C. Stampfer, and H. Missbauer. Comparison of two optimization-based order release models with fixed and variable lead times. *International Journal of Production Economics*, 227:107682, 2020.

D. C. Heath and P. L. Jackson. Modeling the evolution of demand forecasts with application to safety stock analysis in production/distribution systems. *IIE Transactions*, 26(3):17–30, 1994.

R. N. Van Hees and W. Monhemius. *Production and Inventory Control: Theory and Practice*. Red Globe Press, London, 1972.

G. Heisig. Planning stability under (s, S) inventory control rules. *OR Spektrum*, 20: 215–228, 1998.

G. Heisig. Planning stability in material requirements planning systems, 2002.

G. Heisig and M. Fleischmann. Planning stability in a product recovery system. *OR Spektrum*, 23:25–50, 2001.

S. J. Hoekstra and J. H. J. M. Romme. *Op weg naar integrale logistieke structuren*. Wolters-Kluwer, Deventer, 1986.

W. J. Hopp and M. L. Spearman. *Factory Physics*. McGraw-Hill, 2008.

G.J. Van Houtum and W. H. M. Zijm. On the relation between cost and service models for general inventory systems. *Statistica Neerlandica*, 54(2):127–147, 2000.

Y. F. Hung and R. C. Leachman. A production planning methodology for semiconductor manufacturing based on iterative simulation and linear programming calculations. *IEEE Transactions on Semiconductor Manufacturing*, 9(2):257–269, 1996.

Incoterms. Incoterms, 2025. `https://https://en.wikipedia.org/wiki/Incoterms` [Accessed: 10-05-2025].

K. Inderfurth. Nervousness in inventory control: Analytical results. *OR Spektrum*, 16: 113–123, 1994.

K. Inderfurth, A. G. De Kok, and S.D.P. Flapper. Product recovery in stochastic remanufacturing systems with multiple reuse options. *European Journal of Operational Research*, 133(1):130–152, 2001.

D. F. Irdem, N. B. Kacar, and R. Uzsoy. An exploratory analysis of two iterative linear programming-simulation approaches for production planning. *IEEE Transactions on Semiconductor Manufacturing*, 23(3):442–455, 2010.

O. Jabali. *Time and timing in vehicle routing problems*. Phd thesis, Industrial Engineering and Innovation Sciences, 2010.

F. Janssen, A.G. de Kok, and F. van der Duyn Schouten. Approximate analysis of the delivery splitting model. *Journal of the Operational Research Society*, 51(10):1136–1147, 2000. doi: 10.1057/palgrave.jors.2600062. URL `https://doi.org/10.1057/palgrave.jors.2600062`.

T. Jensen. Measuring and improving planning stability of reorder-point lot-sizing policies. *International Journal of Production Economics*, 30-31:167–178, 1993.

T. Jensen. *Planungsstabilität in der Material-Logistik*. Physica, 1996.

N. B. Kacar and R. Uzsoy. Estimating clearing functions for production resources using simulation optimization. *IEEE Transactions on Automation Science and Engineering*, 12(2):539–552, 2015.

N. B. Kacar, D. F. Irdem, and R. Uzsoy. An experimental comparison of production planning using clearing functions and iterative linear programming-simulation algorithms. *IEEE Transactions on Semiconductor Manufacturing*, 25(1):104–117, 2012.

N. B. Kacar, L. M onch, and R. Uzsoy. Modelling cycle times in production planning models for wafer fabrication. *IEEE Transactions on Semiconductor Manufacturing*, 29(2):153–167, 2016.

U. S. Karmarkar. Capacity loading and release planning with work-in-progress (wip) and lead-times. *Journal of Manufacturing and Operations Management*, 2(1):105–123, 1989.

A. Kasper, S. Haeussler, M. Land, and R. Teunter. Better together or divided forever: Integrated or hierarchical material and workload control. *Production and Operations Management*, page in press, 2024.

A. Kefeli and R. Uzsoy. Identifying potential bottlenecks in production systems using dual prices from a mathematical programming model. *International Journal of Production Research*, 54(7):2000–2018, 2016.

O. A. Kilic and S. A. Tarim. An investigation of setup instability in non-stationary stochastic inventory systems. *International Journal of Production Economics*, 133 (1):286–292, 2011.

B. Kim and S. Kim. Extended model for a hybrid production planning approach. *International Journal of Production Economics*, 73:165–173, 2001.

P. Van Laarhoven. Logistiek en supply chains: Innovatieprogramma. Report, Ministerie van Verkeer en Waterstaat, 2008.

A. G. Lagodimos and E. J. Anderson. Optimal positioning of safety stocks in MRP. *International Journal of Production Research*, 31(8):1797–1813, 1993.

K. N. McKay, F. R. Safayeni, and J. A Buzacott. Job-shop scheduling theory: What is relevant? *Interfaces*, 18(4):84–90, 1988.

R. J. Milne, S. Mahapatra, and C. T. Wang. Optimizing planned lead times for enhancing performance of mrp systems. *International Journal of Production Economics*, 167: 220–231, 2015.

H. Missbauer. Order release planning with clearing functions: A queueing-theoretical analysis of the clearing function concept. *International Journal of Production Economics*, 131(1):399–406, 2011.

H. Missbauer. Order release planning by iterative simulation and linear programming: Theoretical foundation and analysis of its shortcomings. *European Journal of Operational Research*, 280:495–507, 2020.

H. Missbauer and R. Uzsoy. *Production Planning with Capacitated Resources and Congestion.* Springer, New York, 2020.

E.G. Negenman. *Material coordination under capacity constraints.* Phd thesis, Industrial Engineering and Innovation Sciences, 2000.

R. S. M. Omar, U. Venkatadri, C. Diallo, and S. Mrishih. A data-driven approach to multi-product production network planning. *International Journal of Production Research*, 55(23):7110–7134, 2017.

J. Orlicky. *Material Requirements Planning.* McGraw-Hill, New York, 1975.

R. Peterson and E. A. Silver. *Decision Systems for Inventory Management and Production Planning.* John Wiley & Sons, New York, 1979.

Y. Pochet and L. Wolsey. *Production Planning by Mixed Integer Programming.* Springer, New York, 2006.

E. Porras and R. Dekker. An efficient optimal solution method for the joint replenishment problem with minimum order quantities. *European Journal of Operational Research*, 174(3):1595–1615, 2006. ISSN 0377-2217. doi: https://doi.org/10.1016/j.ejor.2005.02.056. URL `https://www.sciencedirect.com/science/article/pii/S0377221705002699`.

E. Porras and R. Dekker. A solution method for the joint replenishment problem with correction factor. *International Journal of Production Economics*, 113(2):834–851,

2008. ISSN 0925-5273. doi: https://doi.org/10.1016/j.ijpe.2007.11.008. URL https://www.sciencedirect.com/science/article/pii/S0925527307003386. Special Section on Advanced Modeling and Innovative Design of Supply Chain.

J. Quinonero-Candela, M. Sugiyama, A. Schwaighofer, and N. D. Lawrence, editors. *Dataset Shift in Machine Learning.* MIT Press, 2008.

W. D. Rustenburg, G.-J. van Houtum, and W. H. M. Zijm. Spare parts management for technical systems: resupply of spare parts under limited budgets. *IIE transactions*, 32(10):1013–1026, 2000.

W. D. Rustenburg, G.J. Van Houtum, and W. H. M. Zijm. *Exact and approximate analysis of multi-echelon, multi-indenture spare parts systems with commonality*, pages 143–176. Kluwer, Boston, 2003.

P. Sáez, C. Herrera, and V. Parada. Reducing nervousness in master production planning: A systematic approach incorporating product-driven strategies. *Algorithms*, 16(8): 386, 2023.

C. Schneeweiss. *Distributed Decision Making.* Springer-Verlag, Berlin, 2003.

T. Schoenmeyr and S. C. Graves. Strategic safety stocks in supply chains with evolving forecasts. *Manufacturing & Service Operations Management*, 11(4):657–673, 2009.

T. Schoenmeyr and S. C. Graves. Coordination of multiechelon supply chains using the guaranteed service framework. *Manufacturing & Service Operations Management*, 24(3):1859–1871, 2022.

R. Schonberger. *Japanese manufacturing techniques: Nine hidden lessons in simplicity.* Simon and Schuster, 1982.

B. Selcuk. *Dynamic performance of hierarchical planning systems: modeling and evaluation with dynamic planned lead times.* PhD thesis, EIndhoven University of Technology, 2007.

D. Sever. *Routing in stochastic networks.* Phd thesis, Industrial Engineering and Innovation Sciences, 2014.

C.C. Sherbrooke. Metric: A multi-echelon technique for recoverable item control. *Operations Research*, 16:122–141, 1968.

E. A. Silver, D. F. Pyke, and R. Peterson. *Inventory Management and Production Planning and Scheduling.* John Wiley & Sons, 1998.

J. M. Spitter, A. G. De Kok, and N. P. Dellaert. Timing production in lp models in a rolling schedule. *International Journal of Production Economics*, 93-94:319–329, 2005a.

J.M. Spitter. *Rolling Schedule Approaches for Supply Chain Operations Plannings.* Phd thesis, Industrial Engineering and Innovation Sciences, 2005.

J.M. Spitter, C. A. J. Hurkens, A. G. De Kok, J. K. Lenstra, and E. G. Negenmann. Linear programming models with planned lead times for supply chain operations planning. *European Journal of Operational Research*, 163(3):706–720, 2005b.

A. Srinivasan, M. Carey, and T. E. Morton. Resource pricing and aggregate scheduling in manufacturing systems. Graduate School of Industrial Administration, Carnegie-Mellon University. Pittsburgh, PA, 1988. samizdat.

P. Sutterer, R. Kolisch, and R. Uzsoy. Valuation of hospital resources: An optimization approach using clearing functions. *IISE Transactions on Healthcare Systems Engineering*, 12(4):245–262, 2022.

D. Tas. *Time and reliability in vehicle routing problems.* Phd thesis, Industrial Engineering and Innovation Sciences, 2013.

D. Taş, N. Dellaert, T. Van Woensel, and A. G. De Kok. Vehicle routing problem with stochastic travel times including soft time windows and service costs. *Computers & Operations Research*, 40(1):214–224, 2013.

S.R. Tayur. Computing the optimal policy for capacitated inventory models. *Communications in Statistics-Stochastic Models*, 9:585–598, 1993.

H. C. Tijms. *Stochastic modelling and analysis: a computational approach.* John Wiley & Sons, Chichester, 1986.

M.C. van der Heijden. Supply-rationing in multi-echelon divergent systems. *European Journal of Operational Research*, 101:532–549, 1997.

M.C. van der Heijden, E.B. Diks, and A.G. De Kok. Stock allocation in general multi-echelon distribution systems with (r,s) order-up-to policies. *International Journal of Production Economics*, 49:157–174, 1997.

G.J. van Houtum and W. H. M. Zijm. Computational procedures for stochastic multi-echelon production systems. *International Journal of Production Economics*, 23 (1-3):223–237, 1991.

G.J. Van Houtum, K. Inderfurth, and W. H. M. Zijm. Materials coordination in stochastic multi-echelon systems. *European Journal of Operational Research*, 95:1–23, 1996.

J. H. C. M. Verrijdt and A. G. De Kok. Distribution planning for a divergent depotless two-echelon network under service constraints. *European Journal of Operational Research*, 89(2):341–354, 1996.

M. Vlasiou, I. Adan, and J. Wessels. A lindley-type equation arising from a carousel problem. *Journal of Applied Probability*, 41:1171–1181, 2004.

T. Völker and L. Mönch. Data-driven production planning models for wafer fabs: An exploratory study. *IEEE Transactions on Semiconductor Manufacturing*, 36(3): 1–13, 2023.

T. Vollman, W. Berry, and D. Clay Whybark. *Manufacturing Planning and Control Systems.* Dow Jones-Irwin, Illinois, 1984.

T. E. Vollmann, W. L. Berry, D. C. Whybark, and F. R. Jacobs. *Manufacturing Planning and Control for Supply Chain Management.* McGraw-Hill, New York, 2005.

V. C. S. Wiers and A. G. De Kok. *Designing, selecting, implementing and using aps systems.* Springer, 2017.

O. Wight. *MRP II: Unlocking America's Productivity Potential.* Oliver Wight Limited Publications, Inc., Brattleboro, Vermont, 1981.

J. Wijngaard. The effect of interstage buffer storage on the output of two unreliable production units in series, with different production rates. *AIIE transactions*, 11 (1):42–47, 1979.

T. Ziarnetzky, N. B. Kacar, L. Mönch, and R. Uzsoy. Simulation-based performance assessment of production planning formulations for semiconductor wafer fabrication. In *2015 winter simulation conference (wsc)*, pages 1–10. IEEE, 2015.

T. Ziarnetzky, L. Mönch, and R. Uzsoy. Rolling horizon production planning for wafer fabs with chance constraints and forecast evolution. *International Journal of Production Research*, 56(18):6112–6134, 2018.

T. Ziarnetzky, L. Mönch, and R. Uzsoy. Simulation-based performance assessment of production planning formulations with safety stock and forecast evolution using a large-scale wafer fab model. *IEEE Transactions on Semiconductor Manufacturing*, 33(1):1–12, 2020.

W. H. M. Zijm and A. G. De Kok. Production planning and inventory management in a telecommunication industry. *Engineering Costs and Production Economics*, 15: 267–276, 1989.

P. H. Zipkin. *Foundations of Inventory Management.* Irwin, Burr Ridge, IL, 2000.

G. Zäpfel. Systemanalytische konzeption der produktionsplanung und –steuerung für betriebe der fertigungsindustrie. In C. Zink, editor, *Sozio-technologische Systemgestaltung als Zukunftsaufgabe.* Carl Hanser Verlag, München, 1984. (in German).