Interfaces
Vol. 35, No. 1, January–February 2005, pp. 37–48
ISSN 0092-2102 | EISSN 1526-551X | 05 | 3501 | 0037

INFORMS
DOI 10.1287/inte.1040.0116

Philips Electronics Synchronizes Its Supply Chain to End the Bullwhip Effect

Ton de Kok
Department of Technology Management, Technische Universiteit Eindhoven, Den Dolech 2,
P.O. Box 513, 5600 MB Eindhoven, The Netherlands, a.g.d.kok@tm.tue.nl

Fred Janssen, Jan van Doremalen
CQM, Vonderweg 11, P.O. Box 414, 5600 AK Eindhoven, The Netherlands
{f.b.s.janssen@home.nl, vandoremalen@cqm.nl}

Erik van Wachem, Mathieu Clerkx, Winfried Peeters
Philips Semiconductors, Prof. Holstlaan 4, 5656 AA Eindhoven, The Netherlands
{erik.van.wachem@philips.com, mathieu.clerkx@philips.com, winfried.peeters@philips.com}

Demand variability increases as one moves up a supply chain. The demand for finished products is less variable than for subassemblies, which is less variable than for individual components. This phenomenon is known as the bullwhip or Forrester effect. It increases inventory unnecessarily and makes managing the capacity of equipment and personnel difficult. In 1999, Philips Semiconductors confirmed substantial bullwhip effects in some of its supply chains and began developing a collaborative-planning process and tool to reduce them. It sought to reduce inventory and increase customer-service levels by integrating its supply chain planning and control with those of its customers. By applying stochastic multiechelon inventory theory, it developed an advanced planning and scheduling system that supports weekly collaborative planning of operations by Philips Semiconductors and one of its customers, Philips Optical Storage. The project has brought substantial savings. A conservative estimate shows minimum yearly savings of around US$5 million from $300 million yearly turnover. More important, Philips Optical Storage now has a more flexible and reliable supplier that can virtually guarantee quantities and delivery times. Philips Semiconductor is rolling out its new approach to other customers.

Key words: supply-chain management; collaborative planning; bullwhip effect; multiechelon inventory theory.

"Good evening, Singapore." "Good morning, Eindhoven." "Good morning, Southampton." "Good afternoon, Győr." "Good evening, Shanghai." "Good morning, Eindhoven." "Good evening, Taipei." "Good morning, Eindhoven." On January 23, 2004 the weekly collaborative-planning (CP) meeting between Philips Semiconductors and Philips Optical Storage starts at 10.00 A.M. CET. Present are 14 people from seven locations around the globe in four time zones. The meeting lasts one-and-a-half hours. After quickly checking all inputs (sales plans, work in progress (WIP), and stocks), the CP planning tool calculates the synchronized plan that determines all orders to be released at all links in the supply chain. During the meeting people identify and discuss several problems after which they make some adjustments to the WIP manually and produce a new plan. Regenerating the plan takes only a split second. Nobody pays attention to the fact that a feasible plan has been generated, satisfying several thousand material constraints. They focus on identifying possible follow-up actions to discuss within their organizations or with the subcontractors responsible for PCB assembly. "Any further issues?... No? OK, thanks again. Till next week."

For major international companies the size of Philips Electronics, the complexity of coordinating supply chains is exploding. First, ever-shortening product life cycles combined with long lead times challenge operations and supply chain managers in high-volume electronics. Second, global companies are outsourcing manufacturing and assembly activities, fragmenting the supply chain into operationally and legally independent companies. As a result, companies optimize locally and game the system instead of coordinating and optimizing the entire supply chain. The new competitive battle is no longer between individual companies but between multicompany supply chains. As Charles Fine of MIT remarked in *Clockspeed* (1998), "Competitive advantage is lost or gained by how well a company manages a dynamic web of relationships that run throughout its chain of suppliers, distributors, and alliance partners."

Clearly, firms need new ways of working and software tools to deal with this situation. Because no off-the-shelf packages nor descriptions of best practice were available, we developed a collaborative-planning process and collaborative-planning software.

Organizations Involved

Philips Semiconductors (PS) and Philips Optical Storage (POS) are subsidiaries of Philips Electronics, the world's 10th largest electronics company, with sales of over $30 billion in 2003. Its 164,000 employees in more than 150 countries work in consumer electronics, medical systems, lighting, semiconductors, domestic appliances, and personal care. Its research-and-development budget is over $2.5 billion.

PS, headquartered in Eindhoven, The Netherlands, has over 33,000 employees. With sales of over $5 billion in 2003, the company is one of the world's top semiconductor suppliers. It has 20 manufacturing and assembly sites and a sales organization that delivers in 60 countries.

With 9,000 employees worldwide, POS develops and manufactures optical storage products. Its products include drives, subassemblies, and components for audio, video, data, and gaming playback and rewritable CD and DVD consumer products and PC storage products.

Solving the Bullwhip Effect

Early in 2000, PS began a project to reduce the bullwhip effect in collaboration with its customer POS. The bullwhip is the metaphor for the phenomenon that demand variability increases as one moves up a supply chain. As the furthest upstream link in the high-volume electronics supply chain, PS suffered from this effect. Lee et al. (1997) built upon and extended the ideas of Forrester (1958) to identify common business practices and factors that distort and delay information, such as updating demand forecasts, batching orders, rationing, gaming shortages, and price fluctuations. The collaborative-planning (CP) project was to address these issues through an innovative planning concept and a supporting software system.

The Project

Philips formed a steering committee and a project team consisting of staff and operational people from the two companies, management consultants, and operations research experts. Mathieu Clerkx and Winfried Peeters were members of the steering committee, while Jan van Doremalen, Erik van Wachem, Fred Janssen, and Ton de Kok were members of the project team.

The team started by identifying four key requirements: (1) only through intensive collaboration with partners will our supply chain win; (2) we must share key supply chain information, (3) synchronize decisions on capacities and material flows under high volatility, and (4) decide on supply chain questions very quickly. By introducing a CP process, supported by innovative planning-support software, we believed that we would improve competitiveness by improving customer service, increasing sales and margins, and reducing obsolescence and inventories.

We could expect some resistance from operational users in the two Philips organizations and from the contract manufacturers who had to change their roles in—and impact on—the supply chain planning process. We gained their acceptance by holding workshops with all the primary users and their managers and by conducting intensive face-to-face discussions with all other parties involved. We developed a business case with sound qualitative and quantitative arguments to help us to sell the project to higher management. This task fell to the two champions and coauthors of this paper, Mathieu Clerkx, senior vice president of supply chain management (SCM) and information and communication technology (ICT) at PS, and Winfried Peeters, now with PS but at the time vice president of SCM and ICT at Philips Optical Storage.

Philips' costs for consultants on organizational change and for the development and implementation of the decision-support system were about $1.5 million. It took about two years from initial conception until the operational planning process went live in its final form in January 2002. Between September 2000 and January 2002, we implemented a prototype version of the decision-support system, enabling us to base our development of the software on live experience.

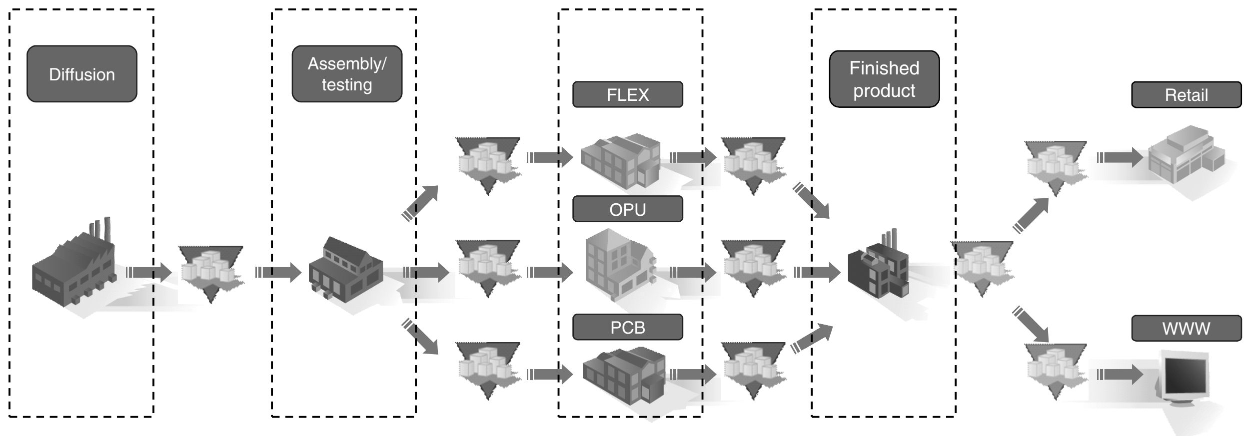

Figure 1: In the value network, the first two blocks cover the IC manufacturing process, the third block represents the DVD drive component manufacturing (printed circuit boards, flex units, and optical pickup units), and the fourth block represents DVD drive assembly.

The Collaborative-Planning-Support Tool

In this paper, we concentrate on the development and use of the CP-support tool. Akkermans et al. (2004) describe the implementation of the CP process in detail.

Ton de Kok suggested that we base the CP-support tool on the synchronized base stock (SBS) policies he had developed to determine capital requirements for safety stock across value networks under demand uncertainty (de Kok and Fransoo 2003). Following this route implied deviating from the dominant paradigm underlying advanced planning and scheduling (APS) systems. The models and solutions in these systems are based largely on mathematical-programming techniques, such as (integer) linear programming, constraint programming, and iterative improvement heuristics (Fleischmann and Meyr 2003). PS hired the Centre for Quantitative Methods (CQM), a consultancy firm employing mostly PhDs in operations research and statisics, to develop the decision-support tool.

The Value Network

To assemble a drive, one needs a number of components and subassemblies, such as an optical pickup unit (OPU), a flex unit, and a printed circuit board (PCB). In turn, to assemble a flex unit or PCB, one needs integrated circuits (ICs), some of which are critical in the sense that they are unique and expensive application-specific ICs (ASICs) and have long manufacturing lead times. The process of manufacturing ICs consists of two major steps: wafer fabrication (including pretest) and IC assembly (including final test) (Figure 1).

PS fabricates the wafers and assembles the ICs. Wafer fabrication is a complex process with long lead times, strong yield variations, and limited capacity. A wafer-fabrication facility costs billions of dollars; such an investment can be justified only for sufficiently high utilization. The fabricated wafers go to assembly centers, which test them for quality and store them. For critical components, the fabrication and testing process takes eight to 12 weeks. The assembly centers assemble and test the ICs and send them to the local industrial warehouses for distribution to the POS's contract manufacturers. Assembly, testing, and transportation take about three weeks. Several contract manufacturers assemble flex units and PCBs. They send the flex units to POS's OPU assembly facilities, which integrate them in OPUs. Typical lead times for assembly and transportation of flex units, PCBs, and OPUs are two weeks. OPUs and PCBs are delivered to POS's drive-assembly centers. The drives then go to regional hubs or customer stock points and finally reach the assembly centers of companies producing DVD players, CD players, and personal computers. Drive assembly and delivery take two to three weeks. So, the total lead time of the value chain is between 17 and 22 weeks.

The Traditional Planning Process

When we began the project, PS and POS had processes in place for high-level coordination (what product portfolio in what numbers) and medium-term planning (what does the market require and what capacity will be made available). We concentrated on the missing link between medium-term planning and execution: short-term planning.

Short-term planning was decentralized and largely disconnected from medium-term planning. The contract manufacturers' and POS assembly plants all had independent weekly planning cycles based on material requirements from the downstream processes in the value chain. They produced netted (modified according to available finished product stocks) plans for material requirements that they sent to the next process upstream in the value chain.

These independent processes caused long information latency (at least six weeks before a change in drive demand downstream affected fabrication decisions upstream) and strong information distortion (caused by poor visibility of material availability, local optimization, and shortage gaming). These weaknesses encouraged all parties to safeguard against uncertainty by creating stocks and, in the process, running obsolescence risks (practices still widespread within the industry). And, in spite of all the safeguarding, deliveries to and from POS were not very reliable.

The Planning Process to Be

The short-term planning problem PS and POS face is a supply-chain-operations-planning problem (de Kok and Fransoo 2003). At the end of each week, the firm must produce a plan prescribing how many of each item in the value network to release to the shop floor. The items to be released are wafers, dies, ICs to be tested, finished ICs at semiconductor warehouses, finished ICs at subcontractor plants, PCBs, and OPUs. The current problem consists of more than 100 items, about 40 of which are end items and about 20, wafers. The horizon of the short-term plan to be decided upon is 26 weeks. A properly coordinated plan should be feasible with respect to both material and resource availability. Philips had no such weekly plan synchronizing order releases along the value network.

Problem Characteristics

The value network from wafer to DVD drive has a general structure: ICs are used in multiple DVD drives, and DVD drives use multiple ICs. Because lead times are so long, planners base most decisions on forecasts. Demand for DVD drives is highly volatile, implying that forecasts of demand in a particular week in the future made on the Wednesday of two consecutive weeks may differ substantially (for example, because of new information on the market or cancellation or acquisition of customer orders). During the project, Philips introduced new products with only rough estimates of total sales during their life cycles. Forecasts of weekly demand up to 26 weeks into the future were therefore inaccurate.

Statistically speaking, one can view demand as stochastic (subject to random fluctuations) and nonstationary (for example, reflecting a product life cycle consisting of an introduction phase, a mature phase, and an end-of-life phase). The throughput times at wafer fabrication and at assembly and testing are stochastic as well: first, because of the economical necessity of using wafer fabs and testing equipment intensively and, second, because of interaction between the items produced on the same resources.

The value network's resource bottleneck is at the most upstream echelon, the wafer fabrication, where the wafer batches dedicated to ICs for POS interact with demand for other wafers. Because of this interaction, proposals from the CP support system to release wafer batches must be checked with resource-consumption information outside the CP support system. Once released, the wafer batch acts as a material constraint for downstream processes while flowing through the supply network. Because the wafer fabrication is the bottleneck, in principle this material constraint is more binding than the resource constraints at downstream processes. Hence, planners manually translate bottleneck-capacity constraints into material constraints, which we take into account in the planning logic.

The yields of wafer fabrication and of assembly and testing are also stochastic and nonstationary, because the introduction of new ICs implies learningcurve effects throughout the value network. As a consequence, the model at the kernel of the decision-support tool had to take into account stochastic nonstationarity of demand, throughput times, and yields. It also had to capture the structural complexity of the value network.

Planning Logic

de Kok and Visschers (1999) formulated the starting point for our development of the planning solution, including algorithms. de Kok and Fransoo (2003) provide an intuitive explanation of this approach. de Kok and Visschers (1999) introduced a class of policies that enable operational control of general multiechelon multi-item systems and computation of close-to-optimal policies within this class under stationary demand (de Kok and Fransoo 2003 later dubbed this class of policies synchronized base stock (SBS) policies). We adapted the planning logic so that we could handle the dynamics sketched above. First, we update the sales plans based on new market information. Second, we update the system state by downloading WIP and stock data from the various enterprise resource planning (ERP) systems. The final set of input data we needed was the planned lead times of work orders to ensure that decisions the CP team made could be executed on the shop floor. The planners determined these lead times to conform to the lead times assumed at the shop-floor level in setting due dates for each order released.

We captured the system's stochastic behavior (the uncertainty with respect to future sales, actual lead times, and yields of work orders) by introducing safety lead times. The planners set these lead times based on their experience. Typically, they set high safety lead times at the start of the product life cycle to protect against potential, yet unpredicted, upsurges in demand. They reduced these safety lead times as soon as they knew more about the product's future demand. In early 2003, Fred Janssen from CQM conducted a study to benchmark their choices for safety lead times against the proposals based on a quantitative SBS-based analysis. As a result of their study, the operational team at PS changed the safety-lead-time parameters.

We combined the information about planned lead times, safety lead times, and sales plans into cumulative demand during cumulative lead times associated with each item modeled. These cumulative demands are equivalent to target order-up-to levels (appendix).

During implementation, we paid a lot of attention to conveying the message that determining the planning logic did not involve an optimization procedure but a calculation procedure. The planning logic originates from stochastic multiechelon inventory theory, which focuses on optimizing planning parameters (in our case, the safety lead times) to cope with uncertainty in demand, lead times, and yields. The planning parameters are considered to be optimal when they yield minimal expected costs over some, possibly infinite, horizon. Once planning parameters are set, the calculation of the order release decisions is more or less trivial (appendix).

The users of the software tool we developed confirmed the effectiveness of the planning logic. Apparently, the solutions derived from a stochastic representation of the real-world problem were consistent with the planners' mental models. Over time the discussion during the weekly CP meetings has changed. Users initially questioned the tools' outcomes. Now they focus on using its interactive problem-solving capabilities.

Problem Solving: The Essence of Speed

A very important contribution of the new software environment is its ability to support problem solving *speedily*. The algorithm can generate feasible plans within seconds. In fact, the calculation of the plan is hardly noticeable to the people participating in the weekly CP meeting. The speed of the algorithm also allows planners to compute multiple plans during the meeting, creating an interactive planning environment.

The software environment also provides strong problem-solving support, used extensively during the CP meetings. One such capability is called backward pegging. It exploits the one-to-one relationship between the shortage of an end item in some future period and a constraining stock on hand or scheduled receipt of one or more upstream items. Thus, the backward-pegging mechanism makes the actual material bottlenecks in the network visible. The users of the system can solve that particular bottleneck manually. For example, they can expedite certain orders or reallocate certain stocks and regenerate the plan to see the consequences of their action.

The Collaborative-Planning Process

The CP process is a weekly process linking monthly supply-chain-capacity agreements, weekly local-production-planning activities, and daily operational execution. The critical outcomes of the process are decisions on (1) the number of wafers to fabricate, (2) the number of ICs to assemble and test, and (3) the number of ICs to ship to the various destinations. The decisions are backed up by sound liability agreements between the involved parties. These agreements guarantee that decisions based on the CP process will be in line with the strategic intent of PS and POS senior management.

The CP process comprises four stages executed in a tightly managed weekly cycle:

Stage 1. Gather data. First, all partners update their parts of the master data in a central database. Then, they collect live data on actual stocks, scheduled receipts, and material in transit and transfer them to the database. They check the data for correctness and consistency.

Stage 2. Decide. Planners make decisions in a virtual meeting with a strict agenda and a knowledgeable moderator. Meetings typically last less than an hour. The partners share their thoughts in a teleconference and interactively view and plan their material flows with the CP planning tool that they share via a net-meeting environment. Typically, the session starts with a review of the previous week's action points and of supply chain parameters, followed by the actual status of the material flows and the new sales forecasts. Using the interactive planning environment, the planners calculate and evaluate alternate scenarios. The planning tool allows for interactive problem solving through its fast algorithm, its transparent planner-friendly solutions, its ability to link downstream supply issues with upstream material-availability problems and vice versa, and its strong user interface.

Stage 3. Escalate. If the planners cannot come to an agreement or decisions fall outside their responsibility area, they refer the issues to the appropriate managers. Remarkable in over 100 weeks of live action, according to the records, planners have not had to call on upper management, whereas before the introduction of this collaboration, upper management involvement was the rule rather than the exception.

Stage 4. Deploy. All decisions are deployed in the organizations involved, that is, PS, POS, and the contract manufacturers.

Business Impact

We established a new planning process and developed and introduced innovative planning-support software. We finished the project successfully in December 2001, and the process has been running ever since. The organizational impact of the CP process and its supporting planning software shows that it is possible to (1) quickly address and intelligently resolve complex material-coordination issues, (2) reveal and capture profitable opportunities, and (3) build trustful relations between independent partners in a supply chain.

In interviews, the key stakeholders, both users and managers, claim that the new way of working and the supporting software environment were instrumental in creating customer intimacy, a sound basis for cooperation and mutual understanding. Their key messages were that the new CP process has greatly reduced information lead time and the CP tool has created greater visibility based on accurate information and enabled decision making based on facts.

The project has affected Philips qualitatively and quantitatively. The obvious and easy-to-verify results include stock reduction (Figure 2), reduction of obsolescence (Figure 3), increased ability to respond to upturns in the market (Figure 4), and improved management of opportunities, such as rerouting available ICs to different DVD drives than initially foreseen (Figure 5). Profit per year measured as percentage of total turnover has increased by 1.5 percent. The yearly benefits in reductions of stock and obsolescence total around US$5 million on a US$300 million turnover.

Detailed evidence of the business impact comes from a detailed analysis of actual operational performance improvements. We gathered real-life data from

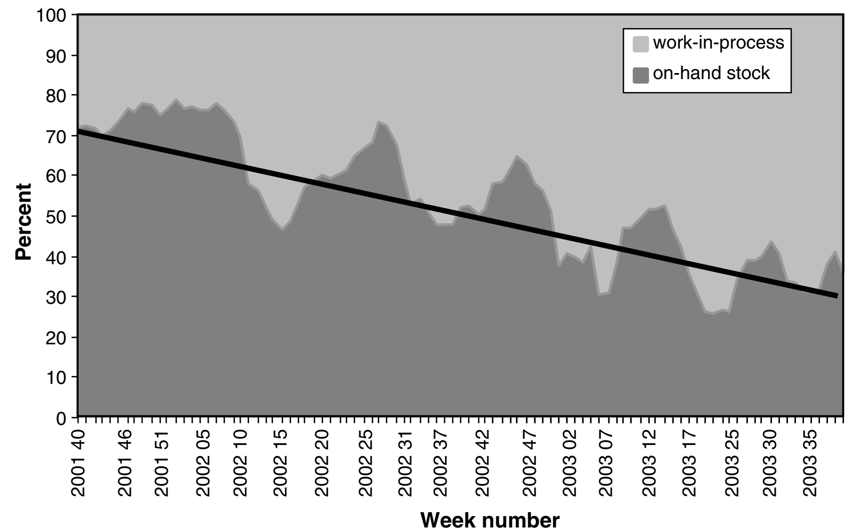

Figure 2: The ratio between on-hand stock and work-in-process improved between 2001 and 2003. (It dropped from 70 percent: 30 percent to 40 percent: 60 percent.)

100 collaborative planning cycles from October 2001 to September 2003. First, the reliability of committed deliveries increased dramatically, because planners had accurate information on what numbers and delivery dates they could and could not offer. Second, our analysis of the data shows a better balance between supply and demand and lower stock levels.

Even more important has been greatly increased flexibility and reliability in serving Philips Optical Storage customers. In the dynamic market of

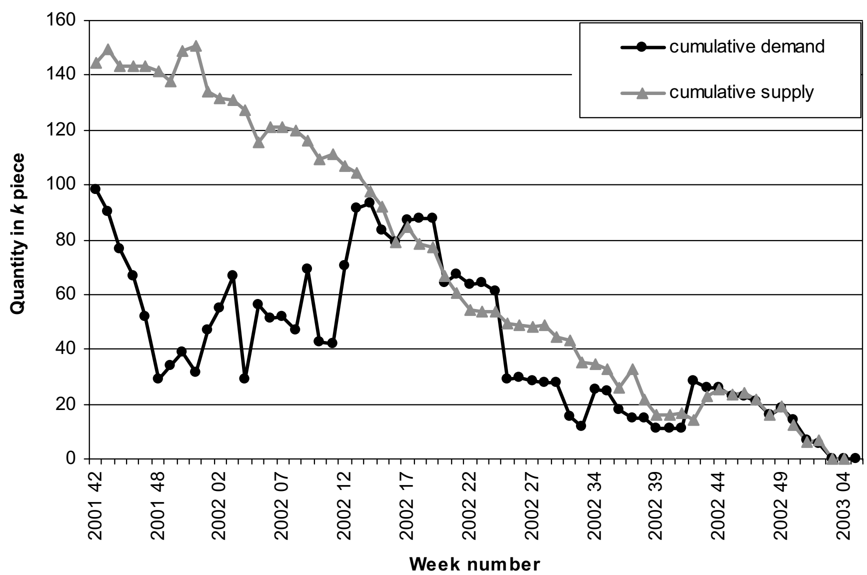

Figure 3: In a successful ramp-down at the end of the life cycle of a product, as a consequence of the new planning process, the gap between demand and supply decreased (first half of the graph), and then supply started following demand closely (second part of the graph) with almost no obsolescence at the end of the life cycle.

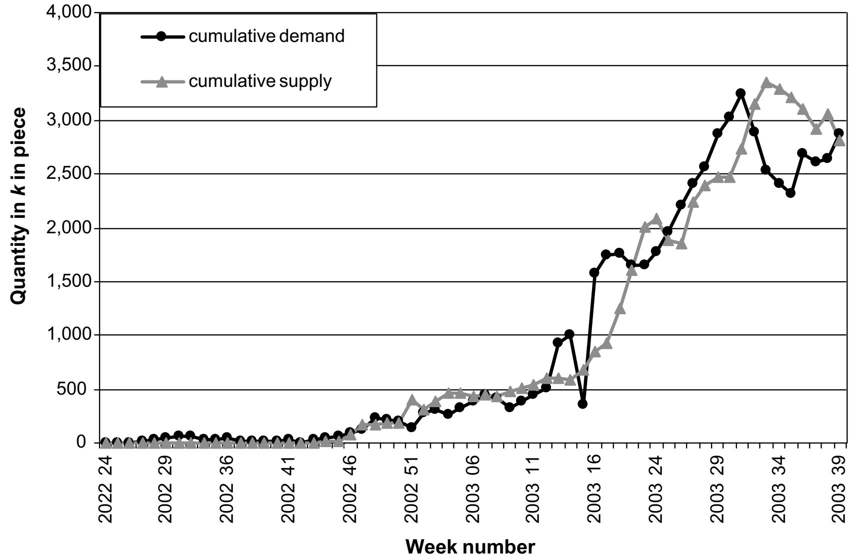

Figure 4: In this successful ramp-up, the supply line closely follows the demand line. Philips Semiconductors had some delivery-lead-time issues before it balanced the chain, but the new behavior was a complete turnaround from previous ramp-ups in which dramatic overshoots or shortages were the rule rather than the exception.

high-tech electronic parts, it is hard to identify and measure all the drivers that increase sales. Instead, Philips Optical Storage must respond rapidly to changes in market conditions.

Top managers testify to the profound impact the new system has had. At PS, Leon Husson, executive vice president of consumer business, and Hein van der Zeeuw, former CEO of POS, say: "The prime

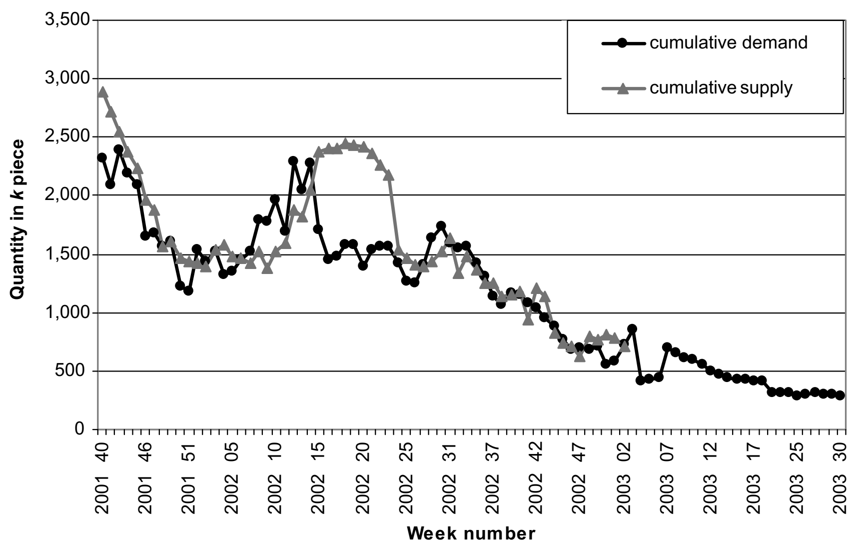

Figure 5: Complete transparency has many advantages. In this example, half way into the life cycle, demand drops very quickly. This drop created a huge obsolescence risk and called for immediate action. Planners discussed the future demand curve in the weekly meeting, and, with a little creative thinking, solved the problem. Philips Semiconductors degraded the integrated circuits to a lower system and sold them for another application, making the best out of this potentially costly problem.

focus of PS is delivering leading-edge systems on silicon. These products have become key components for the company's customers, and the performance of the supply chain network has become a key differentiator. For customers, a well-coordinated supply chain is crucial for fast and flexible responses to changes in the market and a proactive attitude towards market opportunities. The CP project has enabled the required coordination."

According to Leo van Leeuwen, CEO of POS, "the CP process is now well proven and has natural support in the purchasing and supply chain departments. Its contribution to managing effectively the steep ramp-up and ramp-down curves normal in this business is evident to all involved. Currently, use of the tool is being expanded to the company's automotive business line, and discussions are underway for setting up a similar way of working with a solid state laser supplier. This illustrates a high degree of trust in the process: results must be worth the effort to set up such intensive cooperation between business partners." POS champion Winfried Peeters noted that "since the CP support tool went live in January 2002, not a single problem has been escalated to higher management."

Conclusions

CP has been transformed at PS and POS from a buzzword into daily operations. Instead of being involved in myopic fire fighting and escalating operational problems to the executive level, planners are in control and can discuss the tactical and strategic implications of their decisions.

The decision-support tool is the key enabler of this process. Where ERP systems ensure the availability of data about the state of the value network, the CP software synthesizes the data into information focused on synchronized work-order release plans across the value network. The software signals potential problems in filling orders and supports the identification of potentially effective measures.

The weekly CP meeting ensures that planners take such measures with the full support of all the people operationally responsible for order fulfillment in each echelon of the value network.

The underlying logic has proven effective in a highly volatile environment. Although the logic is problem specific, the value-network model is generic. In fact, the value-network model is identical in terms of potential level of detail to the model implicitly assumed in material requirements planning (MRP-I) modules of ERP systems. The fact that the value network model is generic and identical to the MRP-I model makes it transportable to many other situations, whether intracompany or intercompany.

Moreover, the fact that the computational efficiency of the planning logic is comparable to that of MRP-I logic potentially provides an answer to a major gap in current APS suites: the link between short-term planning and shop-floor scheduling.

First and foremost, the implementation has shown the enormous potential of operations research as a discipline and, in particular, the effectiveness of ideas based on insights into complex stochastic problems, such as multi-item multiechelon inventory control. The key to the success of the project has been the intense working relationships between the operations research specialists from the Technische Universiteit Eindhoven and CQM and the operational planners who articulated their needs and critically assessed the plans generated week after week. All this has led to a provably synchronized supply chain from die to DVD and, for Philips, a solution that finally knocks the bullwhip effect on the head.

Appendix. Model and Analysis

We outline the generic model that describes the value network and the algorithm that generates feasible plans for materials for each item in the network. de Kok and Visschers (1999) and de Kok and Fransoo (2003) discuss the underlying ideas from stochastic multiechelon inventory theory supporting the model and the algorithm we refer to.

The Model

We consider an acyclic value network structure with M items. The value network can be described by means of the following sets and their mutual relationships:

- $\mathcal{M}$ set of all items.
- $\mathcal{N}$ set of all end items, i.e., items sold to customers of the value network.
- C_i set of immediate successors or parent set of item i, $i \in \mathcal{M}$.

P_i set of immediate predecessors or child set of item i, $i \in \mathcal{M}$.

F_i set of end items delivered by item i, $i \in \mathcal{M}$.

With each item $i \in \mathcal{M}$, we associate a number of parameters:

a_{ij} number of items i required to produce one unit of item j, $j \in \mathcal{M}$.

L_i lead time of a work order of item i.

L^*_{ij} sum of lead times associated with all items on the path between item i and item j (both items inclusive), $j \in \mathcal{M}$.

ST_i safety lead time associated with item i.

ST^*_{ij} sum of safety lead times associated with all items on the path between item i and item j (both items inclusive), $j \in \mathcal{M}$.

The lead times, L_i, exclude the safety lead times, ST_i. For ease of presentation, we assume that $a_{ij} \in \{0, 1\}$. Clearly, the (safety) lead times associated with items are parameters that have a major impact on the decisions generated. Within the context of this operational planning model, we assume (safety) lead times to be exogenous parameters. The lead times must be derived from actual measurements. Theoretically speaking, the safety lead times may be derived from the analysis of stochastic multiechelon inventory systems (de Kok and Fransoo 2003). In the project, users set the safety lead times used in the planning-software environment.

We assume that, in each review period, a work order is released for all items of the value network. The work-order-release quantities are the decision variables, whose value must be determined. We assume no lot-sizing restrictions for released quantities. We define the following variables for all $i \in \mathcal{M}$ and $t \geq 1$:

$D_i(t)$ forecast of demand for end item i in period t.

$I_i(t)$ net stock of item i at the start of period t.

$IP_i(t)$ inventory position of item i at the start of period t.

$EIP_i(t)$ echelon inventory position of item i at the start of period t.

$SR_i(t)$ scheduled receipt of item i planned to arrive at the start of period t.

$PO_i(t)$ work order of item i released at the start of period t.

Important to the solution sought is that only feasible work orders for materials are released. This implies that the net stock of all nonend items is nonnegative immediately after all orders have been released, i.e.

$$I_i(t) \geq 0, \quad i \in \mathcal{M} \setminus \mathcal{N}, \quad t \geq 1.$$

We derive inventory positions and echelon inventory positions from the following equations:

$$IP_i(t) = I_i(t) + \sum_{s=1}^{L_i-1} SR(t+s), \quad i \in \mathcal{M}, \quad t \geq 1,$$

$$EIP_i(t) = IP_i(t) + \sum_{j \in C_i} EIP_j(t), \quad i \in \mathcal{M}, \quad t \geq 1.$$

In the first of the above equations, we assume that the scheduled receipt arriving at the start of period t is consolidated in $I_i(t)$. We assume that demand forecasts not satisfied from (planned) end-item stocks are backlogged. Without loss of generality, we may assume that we are at the start of period 1, that is, at time 0. At this epoch, the exogenous input to determine the initial state of the system can be described as follows:

$$SR_i(t+s), \quad t = 1, \ldots, L_i - 1, \quad i \in \mathcal{M},$$

$$D_i(t), \quad t = 1, \ldots, T.$$

Despite the fact that in principle only the immediate work-order-release decisions are relevant, below we present an algorithm that generates both immediate work-order-release decisions and planned work-order-release decisions. The latter provide insight into possible future item shortages and overages, which is quite relevant in an environment with short product life cycles and high volatility in demand and supply. Thus, we assume that we have to determine $PO_i(t)$ for all items i until some period T, which as a result is the planning horizon. We assume that T is long enough to accommodate all immediate planning decisions. This implies that

$$T \geq \max_{i,j}(L^*_{ij} + ST^*_{ij}) + 1, \quad i, j \in \mathcal{M}.$$

In what follows, we assume that (planned) events occur in the following order:

(1) Facilities receive scheduled or planned items immediately at the start of a period;

(2) They release work orders for each item immediately after this;

(3) They fulfill customer-demand forecasts and internal work orders just before the end of the period.

Below we describe how to compute the immediate work-order-release decisions. Next, we update the (planned) state of the value network by "executing" these decisions, assuming that scheduled and planned receipts arrive according to their planned lead times and assuming that demand realizations equal demand forecasts. Next, we apply the procedure described below to the new (planned) state of the value network. We repeat this process until the end of the planning horizon. The state-updating procedure can be described as follows:

$$SR_i(t+L_i) = PO_i(t), \quad i \in \mathcal{M}, \quad t \geq 1,$$

$$I_i(t+1) = I_i(t) - D_i(t) + SR_i(t), \quad i \in \mathcal{N}, \quad t \geq 1,$$

$$I_i(t+1) = I_i(t) - \sum_{j \in C_i} a_{ij} PO_j(t) + SR_i(t),$$

$$i \in \mathcal{M} \setminus \mathcal{N}, \quad t \geq 1.$$

Determining Material Feasible Work Orders

Important to a proper sequencing of the decisions to be taken at the start of a particular period t, $t = 1, \ldots, T$, is the acyclic structure of the value network. The computations start with the most upstream items of the value network, that is, items with no predecessors. With a proper low-level coding procedure of the acyclic network, we can recursively determine all subsequent decisions. Our procedure is based on a combination of base-stock policies and linear allocation rules. The procedure yields release decisions that may be inefficient in the sense that the solution proposed may result in unused child items that could be used to satisfy immediate parent-item demand. However, in a multi-item multiechelon system under stochastic demand, this inefficiency need not harm the system's performance, because such residual stock can (should) be used in future periods (de Kok and Visschers 1999). In what follows, we drop the argument t referring to the current period.

Dynamic Base-Stock Levels

The first step in the algorithm is to determine target base-stock levels. Because the echelon stock of an item i indicates the cumulative stock available to cover demand over the lead times $L^*_{i,k} + 1, k \in F_i$, including safety stocks, we define the target base-stock levels S_i as follows:

$$S_i = \sum_{k \in F_i} \left\{ \sum_{s=1}^{L^*_{i,k}+ST^*_{i,k}+1} D_k(s) \right\}, \quad i \in \mathcal{M}.$$

The cumulative safety stock in the echelon of i, SS_i, is defined as

$$SS_i = \sum_{k \in F_i} \left\{ \sum_{s=1}^{L^*_{i,k}+ST^*_{i,k}+1} D_k(s) \right\} - \sum_{k \in F_i} \left\{ \sum_{s=1}^{L^*_{i,k}+1} D_k(s) \right\}, \quad i \in \mathcal{M}.$$

Ordering, Allocation, and Work-Order Release

Standard base-stock policies do not guarantee material feasibility. When an order for an item j derived from the base-stock policy can be satisfied from the available stock of an item $i \in P_j$, then we allocate the required quantity of item i to the order. If this is not the case, because the total required quantity of item i exceeds its available stock I_i, we apply consistent appropriate share (CAS) allocation policies (Van der Heijden et al. 1997) to allocate all available stock.

However, this procedure tackles only the availability issue for item i and all of its successors j. Because we must guarantee that the order from j is feasible for all the children $n \in P_j$, we take PO_j as a minimum of all the quantities allocated from its predecessors. This reasoning can be translated into the following steps:

Let q_j be the unconstrained order from item j, i.e.,

$$q_j = (S_j - EIP_j)^+.$$

Let us consider item $i \in P_j$. We want to determine the quantity $Q_j^{(i)}$, which is the order released for item j if item i would be the only predecessor of item j. Let us assume that for all $m \in C_i$, the unconstrained orders have been determined. Now we can distinguish between two situations,

$$\sum_{m \in C_i} q_m \leq I_i \quad \text{and} \quad \sum_{m \in C_i} q_m > I_i.$$

(1) $\sum_{m \in C_i} q_m \leq I_i$. In this case, we satisfy all orders for item i. Thus, we find $Q_j^{(i)} = q_j$.

(2) $\sum_{m \in C_i} q_m > I_i$. In this case, we must allocate available stock I_i. Van der Heijden et al. (1997) discuss

the consistent appropriate share-rationing policy. The main idea is to allocate the shortage of item i according to the cumulative safety stocks of items $m \in C_i$. Defining EIP_j^+ as the echelon inventory position of j immediately after allocation of the available stock of item i, we find

$$EIP_j^+ = S_j - \frac{SS_j}{\sum_{m\in C_i} SS_m}\left(\sum_{m\in C_i} q_m - I_i\right).$$

It is well known that due to imbalance, it may be possible that $EIP_j^+ < EIP_j$, which implies that a negative quantity would be allocated to item j. Taking this into account, we calculate $Q_j^{(i)}$ in this case as

$$Q_j^{(i)} = \frac{\max(0, EIP_j^+ - EIP_j)}{\sum_{m\in C_i} \max(0, EIP_m^+ - EIP_m)} I_i.$$

Thus, we have found $Q_j^{(i)}$ for the two possible cases. Then, we determine the order released for item j from

$$PO_j = \min_{n\in P_j} Q_j^{(n)}.$$

This concludes the algorithm.

References

Akkermans, H. A., P. Bogerd, J. B. M. van Doremalen. 2004. Travail, transparency and trust: A case study of computer-supported collaborative supply chain planning in high-tech electronics. *Eur. J. Oper. Res.* **153**(2) 445–456.

de Kok, A. G., J. C. Fransoo. 2003. Planning supply chain operations: Definition and comparison of planning concepts. A. G. de Kok, S. C. Graves, eds. *Handbooks in Operations Research and Management Science*, *Vol.* 11. *Supply Chain Management: Design, Coordination and Operation*, Ch. 12. North-Holland, Amsterdam, The Netherlands.

de Kok, A. G., J. W. C. H. Visschers. 1999. Analysis of assembly systems with service level constraints. *Internat. J. Production Econom.* **59**(1–3) 313–326.

Fine, C. H. 1998. *Clockspeed: Winning Industry Control in the Age of Contemporary Advantage.* Perseus Books, Reading, MA.

Fleischmann, B., H. Meyr. 2003. Planning hierarchy, modeling, and advanced planning systems. A. G. de Kok, S. C. Graves, eds. *Handbooks in Operations Research and Management Science*, *Vol.* 11. *Supply Chain Management: Design, Coordination and Operation*, Ch. 9. North-Holland, Amsterdam, The Netherlands.

Forrester, J. W. 1958. System dynamics: A major breakthrough for decision makers. *Harvard Bus. Rev.* **36**(4) 37–66.

Lee, H. L., V. Padmanabhan, S. Whang. 1997. The bullwhip effect in supply chains. *Sloan Management Rev.* **38**(3) 93–102.

Van der Heijden, M. C., E. B. Diks, A. G. de Kok. 1997. Stock allocation in general multi-echelon distribution systems with (*R*,*S*) order-up-to-policies. *Internat. J. Production Econom.* **49**(2) 157–174.