

Optimizing the Supply Chain Configuration for New Products

Stephen C. Graves

Leaders for Manufacturing Program and Sloan School of Management, Massachusetts Institute of Technology,
Cambridge, Massachusetts 02139-4307, sgraves@mit.edu

Sean P. Willems

School of Management, Boston University, 595 Commonwealth Avenue, Boston, Massachusetts 02215,
willems@bu.edu

We address how to configure the supply chain for a new product for which the design has already been decided. The central question is to determine what suppliers, parts, processes, and transportation modes to select at each stage in the supply chain. There might be multiple options to supply a raw material, to manufacture or assemble the product, and to transport the product to the customer. Each of these options is differentiated by its lead time and direct cost added. Given these various choices along the supply chain, the configuration problem is to select the options that minimize the total supply chain cost. We develop a dynamic program with two state variables to solve the supply chain configuration problem for supply chains that are modeled as spanning trees. We illustrate the problem and its solution with an industrial example. We use the example to show the benefit from optimization relative to heuristics and to form hypotheses concerning the structure of optimal supply chain configurations. We conduct a computational experiment to test these hypotheses.

Key words: dynamic programming application; multiechelon inventory system; multistage supply chain application; total landed-cost optimization; supply chain sourcing and design; supply chain configuration for new products

History: Accepted by Fangruo Chen, supply chain management; received March 25, 2002. This paper was with the authors 1 year and 3 months for 2 revisions.

1. Introduction

In this paper, we examine optimal configuration strategies for new product supply chains. Our intent is to develop a decision-support tool that product managers can use during the product development process when the product's design has been decided but the vendors, manufacturing technologies, and shipment options have not yet been determined. Our supply chain design framework considers three specific costs that are relevant when configuring new supply chains: cost of goods sold and the inventory holding costs for both safety stock and pipeline stock. The supply chain configuration problem chooses a sourcing option for each stage of the supply chain so as to minimize the sum of these costs.

For most structured product development processes, there is a milestone when the materials management organization (MMO) sources the new product's supply chain. The product's functionality has already been determined at this point. There are several available sourcing options at each stage. Examples include multiple vendors to supply a raw material, several manufacturers or technologies to

manufacture or assemble the product, and numerous transportation modes to deliver the finished product.

The role of the MMO is to identify the options that can satisfy each function and then to decide which options to select. Options differ in terms of their direct costs and lead times. Therefore, choices in one portion of the supply chain can affect the costs and responsiveness of the rest of the supply chain. The tradeoff facing the MMO is whether to create a higher unit manufacturing cost with a more responsive supply chain versus a lower manufacturing cost with a less responsive supply chain.

This problem integrates and builds upon ideas from the research literature in the areas of multiechelon inventory theory and network design.

From the literature on multiechelon inventory, numerous papers address optimizing safety stock placement across the supply chain, for example, Ettl et al. (2000) and Graves and Willems (2000a). These papers, and many of their cited references, optimize safety stock levels for an established supply chain. Because these models consider existing supply chains, there is already one option chosen at each stage.

Therefore, the cost of goods sold (COGS) is set, as is the holding cost of the pipeline inventory, and these costs do not enter into the analysis.

Inderfurth (1993) jointly considers the optimization of safety stock costs and lead times for a single production process producing multiple end items. The optimization captures the impact that the finished goods' lead times have on the safety stock in the supply chain. However, the model only considers changing the configuration at one stage in the supply chain and only considers safety stock costs.

Research on network design focuses on developing the optimal manufacturing and distribution network for a company's entire product line. Geoffrion and Powers (1995) describe the evolution and primary assumptions in this field. These approaches generally formulate large-scale integer linear programs that capture the relevant fixed and variable operating costs for each facility; these variable costs go beyond manufacturing costs to include tariffs and taxes. This stream of research (for instance, Geoffrion and Graves 1974, Arntzen et al. 1995) differs from this paper in its level of detail and scope. Network design focuses on the design of two or three echelons in the supply chain for multiple products; the supply chain configuration problem focuses on a single product family at the supply chain level, allowing it to model all the echelons in the supply chain and to explicitly capture the impact of variability on the supply chain.

This paper is structured as follows. In §2, we present the notation and assumptions. We formulate in §3 the optimization problem and demonstrate how to solve the problem by dynamic programming in §4. In §5, we provide a case study showing the benefits of optimizing the supply chain configuration versus two heuristic approaches: choosing the cheapest unit cost supply chain and choosing the most responsive supply chain. We conclude this section with a set of observations regarding the structure of optimal supply chain configurations, and in §6 present a design of experiments to validate these observations. The conclusion and next steps are in §7.

2. Notation and Assumptions

The modeling framework follows that of Graves and Willems (2000a). We will not repeat the discussion justifying the assumptions when they are the same as in the earlier paper.

2.1. Option Definition

We model a supply chain as a network of stages, where each stage represents a necessary function, such as procurement, assembly, or transportation. For each stage, one or more options exist that can satisfy the stage's functional requirement. For example, if a stage represents the procurement of a metal housing,

then one option might be a locally based high-cost provider and another option could be a low-cost international supplier. For each stage, we will select a single option. We characterize an option at a stage by its direct cost added and lead time. When a stage reorders, the lead time is the time to perform the function at the stage, provided all the inputs are available; we assume that lead times are deterministic. An option's direct cost represents the direct material and direct labor costs associated with the option. If the option were the procurement of a raw material from a vendor, then the direct cost is the purchase price plus any transportation and labor costs to unpack and inspect the product.

2.2. Periodic-Review Base-Stock Replenishment Policy

We assume that each stage operates according to a periodic review policy with a common review period. Each period each stage observes demand either from external customers or from its downstream stages, and places an order on its suppliers to replenish the observed demand. In effect, each stage operates with a one-for-one or base-stock replenishment policy. We assume no time delay in ordering so that in each period all stages see the external customer demands. We can extend the model to permit a deterministic delay when passing the order information up the supply chain. For determining the safety stocks, this information delay is equivalent to increasing the production lead time by the delay; however, the information delay does not affect the pipeline stock.

2.3. Demand Process

We assume that external demand occurs only at nodes that have no successors, which we term demand nodes or stages. For each demand node j , we assume that the end-item demand comes from a stationary process, with average demand per period μ_j .

An internal stage has only internal customers or successors; its demand in period t is the sum of the orders placed by its immediate successors. Because each stage orders according to a base-stock policy, the demand at internal stage i is

$$d_i(t) = \sum_{(i,j) \in A} d_j(t),$$

where $d_j(t)$ denotes the demand at stage j in period t , A is the arc set for the network representation of the supply chain. For each notation, and without loss of generality, we assume that if there is an arc from i to j , then one unit of i is required to produce one unit of j . The average demand rate for stage i is

$$\mu_i = \sum_{(i,j) \in A} \mu_j.$$

We assume that demand at each stage j is *bounded* by the function $D_j(\tau)$, where, for any period t and for

any τ , we have

$$D_j(\tau) \geq d_j(t - \tau + 1) + d_j(t - \tau + 2) + \dots + d_j(t).$$

We define $D_j(0) = 0$ and assume that $D_j(\tau)$ is increasing and concave.

2.4. Guaranteed Service Times

We assume that each stage j promises a guaranteed outbound service time s_j^{out} by which the stage will satisfy its demand, either from internal or external customers. That is, the customer demand at time t , $d_j(t)$, must be filled by time $t + s_j^{\text{out}}$. Furthermore, we assume that stage j provides 100% service for the specified service time: Stage j delivers exactly $d_j(t)$ to the customer at time $t + s_j^{\text{out}}$. Demand nodes also receive as input a maximum service time, S_j , which will constrain the choice of s_j^{out} .

We assume that an internal stage i quotes the same outbound service time, s_i^{out} , to all of its downstream customers; Graves and Willems (2000a) describe how to extend the model to permit customer-specific service times.

We define s_i^{in} to be the inbound service time for stage i . For stages with one or more upstream adjacent stages, the inbound service time for stage i equals the maximum of the service times quoted to stage i by its suppliers. Thus, when stage i reorders, the time to receive all the required inputs from the suppliers is s_i^{in} .

3. Optimization Model

We formulate the supply chain configuration problem as an optimization problem for which the decision variables are the options and the service times:

$$\mathbf{P} \quad \min \sum_{i=1}^N \left[\alpha c_i [D_i(s_i^{\text{in}} + t_i - s_i^{\text{out}}) - (s_i^{\text{in}} + t_i - s_i^{\text{out}})\mu_i] + \alpha \left(c_i - \frac{x_i}{2} \right) t_i \mu_i + \beta x_i \mu_i \right]$$

such that

$$\sum_{k=1}^{O_i} T_{ik} y_{ik} - t_i = 0 \quad \text{for } i = 1, 2, \dots, N, \quad (1)$$

$$\sum_{k=1}^{O_i} C_{ik} y_{ik} - x_i = 0 \quad \text{for } i = 1, 2, \dots, N, \quad (2)$$

$$c_i - \sum_{j:(j,i) \in A} c_j - x_i = 0 \quad \text{for } i = 1, 2, \dots, N, \quad (3)$$

$$s_i^{\text{in}} \geq s_j^{\text{out}} \quad \text{for } i = 1, 2, \dots, N, j: (j, i) \in A, \quad (4)$$

$$s_i^{\text{in}} + t_i - s_i^{\text{out}} \geq 0 \quad \text{for } i = 1, 2, \dots, N, \quad (5)$$

$$s_j^{\text{out}} \leq S_j \quad \text{for all demand nodes } j, \quad (6)$$

$$s_i^{\text{in}}, s_i^{\text{out}} \geq 0 \text{ and integer} \quad \text{for } i = 1, 2, \dots, N, \quad (7)$$

$$\sum_{k=1}^{O_i} y_{ik} = 1 \quad \text{for } i = 1, 2, \dots, N, \quad (8)$$

$$y_{ik} \in \{0, 1\} \quad \text{for } i = 1, 2, \dots, N, 1 \leq k \leq O_i, \quad (9)$$

where

O_i = number of candidate options at stage i ,

C_{ik} = direct cost added of the k th option at stage i ,

T_{ik} = lead time of the k th option at stage i ,

$D_i(\cdot)$ = maximum demand function for stage i ,

α = scalar representing the holding cost rate,

β = scalar converting the model's underlying time unit into the company's time interval of interest,

μ_i = mean demand rate at stage i ,

c_i = cumulative cost at stage i ,

t_i = selected option's lead time at stage i ,

x_i = selected option's cost at stage i , and

y_{ik} = indicator variable, which equals 1 if stage i 's k th option is selected and 0 otherwise.

The objective function has three terms, each corresponding to a component of the supply chain configuration cost. The first term represents stage i 's safety stock cost, which is a function of the stage's net replenishment time and demand characterization. The holding cost at stage i equals the cumulative cost of the product at stage i times the holding cost rate. The second term expresses the pipeline stock cost as the product of the holding cost rate, the average cost of the product at the stage, and the expected amount of pipeline stock. The third term, COGS, represents the total cost of all the units that are delivered to customers during a company-defined interval of time. The incremental contribution to COGS is calculated at each stage by a product of the average demand at the stage, the option's cost, and a scalar β , which expresses COGS in the same units as pipeline and safety stock cost.

Constraints (1) and (2) define the cost and lead time for each stage, as they depend on the option chosen. Constraint (3) calculates the cumulative cost at each stage. Constraints (4)–(7) assure that the service times are feasible. In particular, the inbound service time at every stage is at least as large as the largest outbound service time quoted to the stage; the net replenishment time of each stage is nonnegative; the outbound service times to the customer must be no greater than the user-defined maximums; and service times must be nonnegative and integer. The last two constraints, (8) and (9), enforce the sole sourcing of options.

We refer the reader to Willems (1999) for more discussion and for details regarding the underlying structure and formulation of the mathematical program.

4. Dynamic Programming Formulation

In this section, we describe how to solve **P** by dynamic programming when the underlying network is a spanning tree. We solve **P** by decomposing the problem into N stages, where N is the number of nodes in the spanning tree and we can number the nodes such that for each node i there is at most one adjacent node with a label higher than i (see Graves and Willems 2000a). We term the adjacent node with a higher label to be the parent node, which we denote by $p(i)$ for node i ; the parent can be upstream or downstream of node i . Furthermore, we define N_i to be the subset of nodes $\{1, 2, \dots, i\}$ that are connected to i on the subgraph with node set $\{1, 2, \dots, i\}$. We can determine N_i by the following equation:

$$N_i = \{i\} + \bigcup_{h < i, (h, i) \in A} N_h + \bigcup_{j < i, (i, j) \in A} N_j.$$

The dynamic program evaluates a functional equation for each node, proceeding sequentially from 1 to N . The solution at each node is the solution to **P** for the subgraph N_i . Intuitively, as the dynamic program progresses through the network, it is creating a series of solved subnetworks that grow and combine until eventually the entire network is solved. After each iteration, the current node's parent is the only node adjacent to the subnetwork that is unsolved.

We first formulate the functional equations for the dynamic programming recursions and then present the dynamic program.

4.1. Functional Equation Development

There are two forms of the functional equation, depending on the node's orientation in the network. First, when the parent of node i is downstream from node i , we define the function $f_i(c^T, s^{\text{out}})$ as the minimum cost for the supply chain configuration in a subnetwork with node set N_i , where the cumulative unit cost at stage i is c^T and the outbound service time is s^{out} . Second, when the parent of node i is upstream from node i , we define the function $g_i(c^1, s^{\text{in}})$ as the minimum cost for the supply chain configuration in a subnetwork with node set N_i , where stage i 's parent has a cumulative unit cost c^1 and quotes an outbound service time of s^{in} to stage i .

To develop the functional equations, we first define $z_{ik}(s^{\text{in}}, c^1, c^2, s^{\text{out}})$ as the supply chain cost for the subnetwork with node set N_i when option k is selected for stage i . $z_{ik}(\cdot)$ is a function of four arguments: s^{in} is the inbound service time to stage i ; c^1 is the cumulative cost of the parent if the parent is upstream of stage i ; c^2 is the cumulative cost for all of the other upstream adjacent stages; and s^{out} is stage i 's

outgoing service time.

$$\begin{aligned} z_{ik}(s^{\text{in}}, c^1, c^2, s^{\text{out}}) &= \alpha c^T [D_i(s^{\text{in}} + T_{ik} - s^{\text{out}}) - (s^{\text{in}} + T_{ik} - s^{\text{out}})\mu_i] \\ &+ \alpha \left(c^T - \frac{C_{ik}}{2} \right) T_{ik} \mu_i + \beta C_{ik} \mu_i + \sum_{\{j: (i, j) \in A, j < i\}} g_j(c^T, s^{\text{out}}) \\ &+ \min_{\sum_{\{h: (h, i) \in A, h < i\}} c_h = c^2} \left\{ \sum_{\{h: (h, i) \in A, h < i\}} f_h(c_h, s^{\text{in}}) \right\}, \end{aligned} \quad (10)$$

where $c^T = c^1 + c^2 + C_{ik}$. The first three terms correspond to the safety stock cost, direct manufacturing cost, and pipeline stock cost at stage i . These costs are a function of the option selected at the stage, the service times, and the total incoming cost to stage i . The incoming cost to stage i is the sum of two quantities: the parent's cumulative cost and the cumulative cost from all other upstream adjacent stages. These quantities would be zero if there were no upstream stages.

The fourth term corresponds to the nodes in N_i that are downstream from node i . For each node j that is a customer to node i , we include the minimum supply chain cost at stage j as a function of stage i 's contribution to the cumulative cost at stage j and the service time i quotes j .

The fifth term corresponds to the nodes in N_i that are upstream from i . This term consists of the minimum supply chain cost for the configuration upstream from stage i that is capable of producing a cumulative cost c^2 . The incoming service time to stage i (s^{in}) is the maximum service time that is being quoted to stage i . Therefore, s^{in} is an upper bound on the outbound service time that each upstream stage can quote.

We evaluate $f_i(c^T, s^{\text{out}})$ when the parent is downstream from the current node, as follows:

$$\begin{aligned} f_i(c^T, s^{\text{out}}) &= \min_{k, s^{\text{in}}} \{ z_{ik}(s^{\text{in}}, 0, c^T - C_{ik}, s^{\text{out}}) \} \\ \text{s.t. } &\max(0, s^{\text{out}} - T_{ik}) \leq s^{\text{in}} \leq M_i - T_{ik}, \\ &1 \leq k \leq O_i, \quad s^{\text{in}} \geq 0 \text{ and integer,} \end{aligned}$$

where M_i is the maximum possible replenishment time for node i , defined as $M_i = \max\{T_{ik} \mid k: 1 \leq k \leq O_i\} + \max\{M_h \mid h: (h, i) \in A\}$. The lower bound on s^{in} comes from **P**, and its upper bound is by the definition of M_i . The minimization is over the feasible set of options at the stage. By definition, $f_i(c^T, s^{\text{out}})$ is solved when $(i, p(i)) \in A$. Therefore, the sum of the direct costs from all upstream stages must equal $c^T - C_{ik}$, and c^1 is zero in the specification of $z_{ik}(\cdot)$.

When the parent is upstream from the current node, we evaluate the functional equation, $g_i(c^1, s^{\text{in}})$, as

follows:

$$g_i(c^1, s^{\text{in}}) = \min_{k, c^2, s^{\text{out}}} \{z_{ik}(s^{\text{in}}, c^1, c^2, s^{\text{out}})\}$$

$$\text{s.t. } 0 \leq s^{\text{out}} \leq s^{\text{in}} + T_{ik},$$

$$1 \leq k \leq O_i, \quad s^{\text{out}} \geq 0 \text{ and integer.}$$

The range for c^2 is bounded and the bounds can easily be established from the cost parameters in the network. Because $g_i(c^1, s^{\text{in}})$ is solved when $(p(i), i) \in A$, then c^1 and c^2 can both be nonzero. If node i is a demand node, we have the additional constraint that s^{out} cannot exceed its maximum service time, i.e., $s^{\text{out}} \leq S_i$. Again, the minimization can be done by enumeration.

4.2. Dynamic Program

The dynamic programming algorithm is now as follows:

1. For $i := 1$ to $N - 1$
 - 1a. If $p(i)$ is downstream from i , evaluate $f_i(c^T, s^{\text{out}})$ for $s^{\text{out}} = 0, 1, \dots, M_i$, $c^T \in X_i$, where X_i is the set of feasible cumulative costs at stage i .
 - 1b. If $p(i)$ is upstream from i , evaluate $g_i(c^1, s^{\text{in}})$ for $s^{\text{in}} = 0, 1, \dots, M_i$, $c^1 \in X_{p(i)}$, where $X_{p(i)}$ is the set of feasible cumulative costs at stage i 's parent node.
2. For $i := N$ evaluate $g_i(0, s^{\text{in}})$ for $s^{\text{in}} = 0, 1, \dots, M_i$.
3. Minimize $g_N(0, s^{\text{in}})$ for $s^{\text{in}} = 0, 1, \dots, M_N$ to obtain the optimal objective function value. The optimal set of service times and options are found by the standard backtracking procedure for a dynamic program.

The computational complexity of the algorithm is of order $k^N N M^2$, where k is the maximum number of options at any stage, N is the number of nodes, and M is the maximum replenishment time in the network, which is bounded by the sum of the longest lead time at each stage. For each of the N nodes in the network there are at most k^N functional equations to evaluate over the entire range of inbound or outbound service times. We implemented the algorithm in C++ programming language. The run times for real problems with 75–100 nodes and two options per node are effectively instantaneous on a Pentium PC with a 1.7 gigahertz Intel processor.

5. Application: Delayed Differentiation in Computer Manufacturing

This section presents the results from a four-week diagnostic exercise that was conducted at a Fortune 100 computer manufacturer.¹

¹The data in this section have been disguised to protect the company's proprietary information. The insights drawn from the disguised data are the same insights that were drawn from the real data.

5.1. Current Process Description

The company employs a target costing approach (Ansari and Bell 1997) when designing new product supply chains. In brief, the market price for the product is set from outside the product design group. Two common reasons for this are (1) that the product faces many competitors and the firm will be a price taker and (2) that another department within the company, for example, marketing, specifies the product's selling price. Next, a gross margin for the product is specified, typically by senior management or corporate finance. The combination of the prespecified selling price and the gross margin target dictates the product's maximum unit cost.

The product's unit manufacturing cost (UMC) is the sum of the direct costs for the production of a single unit of product. Typical costs include raw materials, the processing cost at each stage, and transportation. The maximum unit cost acts as an overall budget for the product's UMC.

From an organizational perspective, the supply chain development core team is composed of an early supply chain enabler and one or two representatives associated with each of the product's major subassemblies. The early supply chain enabler is responsible for shepherding the product through the product development process. This individual is brought in during the early design phase and will stay with the project until it achieves volume production.

The core team will allocate the UMC budget across the major subassemblies. This is not an arbitrary process. The team will rely on a number of factors, including competitive analysis, past product history, future cost estimates, and value engineering. Once the subassembly budgets are set, the design teams for each subassembly are charged with producing a subassembly that can provide the functionality required subject to its budget constraint. Even if these groups incorporate multidisciplinary teams and concurrent engineering, they will still operate within their own budget constraints.

In much the same way that the UMC is allocated to the subassemblies, each subassembly group allocates its budget and decides what processes and components to use. There are numerous factors to consider when sourcing a component, some of which include functionality, price, vendor delivery history, vendor quality, and vendor flexibility. Because many of these factors are difficult to quantify, the team establishes a minimum threshold for each of the intangible factors. Suppliers that meet or exceed the thresholds are considered as candidates or options.

The company's current practice is essentially to choose the lowest unit cost option from the set of options that satisfy the intangible factors. In the framework of the supply chain configuration problem, this corresponds to choosing the option with

the lowest cost added at each stage, regardless of its lead time. While this is admittedly a heuristic, there are several reasons the company does this. First, as mentioned earlier, all other factors besides cost are difficult, if not impossible, to quantify. For example, the company only wants to do business with suppliers that have been certified. The certification process involves a rigorous review of the supplier’s quality practices, but given two certified suppliers, there is no mechanism to view one supplier’s quality as superior to the other. Second, the UMC of the product will dictate whether the business case to launch the product is successful. If the UMC were not low enough to meet the gross margin target, then the project will be terminated. Therefore, there is tremendous pressure to focus on the UMC at the expense of other considerations. Finally, the team that designs the supply chain is not the same team that has to manage the supply chain. Although choosing parts with long lead times might significantly increase the supply chain’s inventory requirements, this dynamic has not been explicitly considered during the new product’s business case analysis.

5.2. Notebook Computer Case Study

A notebook computer consists of three major sub-assemblies: the liquid crystal display (LCD), the circuit boards, and the housing. The LCD is a standard component that is purchased from an external vendor. The housing is a custom-designed product that is also sourced from an external vendor.

To create the circuit boards, components are purchased from external vendors and assembled by a

contract manufacturer. The assembly process involves assembly of the components and quality testing and creates a generic notebook computer. The generic notebook is then customized with either a gray or blue cover. The standard gray variant serves two different markets: U.S. demand and export demand. The blue variant is a new introduction for the U.S. market.

A graphic depiction of the supply chain is shown in Figure 1.

The circuit board assembly is depicted at the top left of the figure. The components for the circuit board are grouped according to their traditional procurement lead times. The LCD display and metal housing are procured from outside suppliers. The battery is included separately from the other miscellaneous components because it is an expensive accessory with an especially long procurement lead time. After assembly, the generic notebook is customized with either a blue or gray cover and shipped to the appropriate demand location.

Table 1 contains the options available when sourcing this supply chain. The company operates on a five-day work week and there are 250 work days in the year. The annual holding cost rate is 45%.

For each stage, option 1 reflects the option that was chosen for the existing supply chain. The additional options were judged by the materials management group to be alternatives that were feasible options that were not selected.

For the circuit board’s raw materials, the different options refer to different classes of service that the distributor is willing to provide. The head of materials management for the electronics subassembly

Figure 1 Notebook Computer Supply Chain

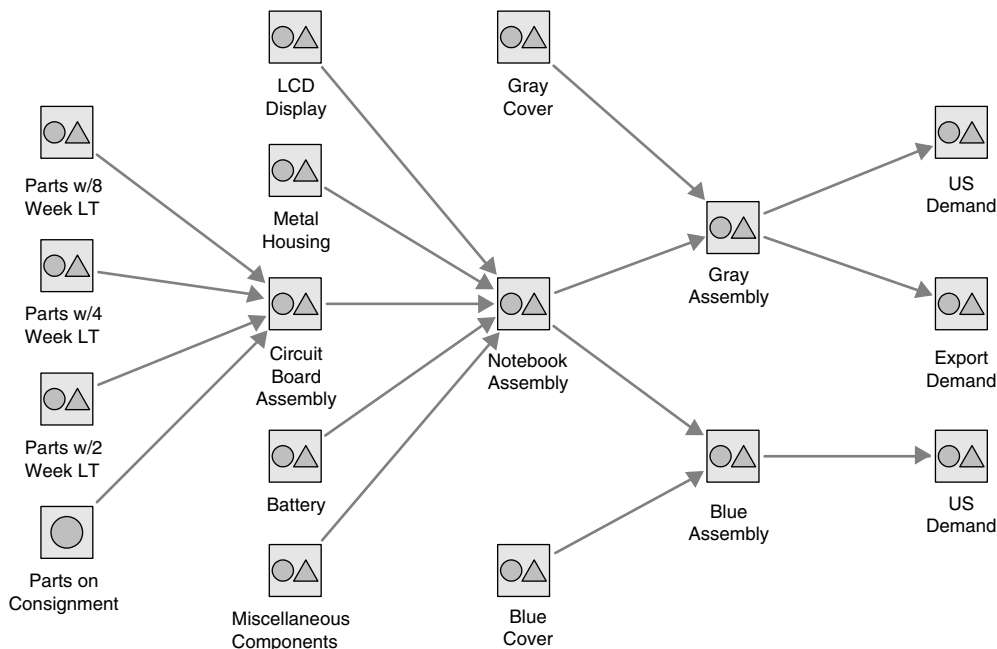


Table 1 Options for Notebook Computer

Component/process description	Option	Lead time	Cost added (\$)
Parts w/ eight-week lead time	1	40	130.00
	2	20	133.25
	3	10	134.91
	4	0	136.59
Parts w/ four-week lead time	1	20	200.00
	2	10	202.50
	3	0	205.03
Parts w/ two-week lead time	1	10	155.00
	2	0	156.93
Parts on consignment	1	0	200.00
Circuit board assembly	1	20	120.00
	2	5	150.00
LCD display	1	60	300.00
	2	5	350.00
Miscellaneous components	1	30	200.00
Metal housing	1	70	225.00
	2	30	240.00
Battery	1	60	40.00
	2	20	45.00
Notebook assembly	1	5	120.00
	2	2	132.00
Gray cover	1	40	5.00
	2	15	5.50
Blue cover	1	40	5.00
	2	15	5.50
Gray assembly	1	1	30.00
Blue assembly	1	1	30.00
U.S. demand—gray	1	5	12.00
	2	1	20.00
Export demand—gray	1	15	15.00
	2	2	30.00
U.S. demand—blue	1	5	12.00
	2	1	20.00

Notes. The example presented here is a slight simplification of the actual diagnostic project. The original diagnostic modeled a 34-stage supply chain with 72 options. The main sources of simplification are the modeling of only two end items versus five end items and the merging of multiple miscellaneous components into a single stage; merged stages include cables, disk drive, and keyboard. The simplification does not change the underlying results, nor does it change the recommendations made.

estimated that converting an eight-week lead time part to a consignment part would increase the purchase price by 5%. We used this information to estimate the cost of reducing one week of lead time for each electronic part as 0.625% of the part’s selling price. In effect, the increase in cost represents the cost incurred by the distributor to inventory the part.

For other components like the metal housing, LCD display, covers, and battery, the costs reflect either different providers or different terms from the same vendor; these contract terms were specified in the request for quotation (RFQ) process.

As a rule of thumb, the company valued one hour of processing time at \$60. Recall that the definition of lead time includes the waiting time at a stage plus the

actual processing time at the stage. Therefore, a slight increase in the available capacity at a stage can dramatically reduce the stage’s lead time. For example, by adding \$30 to the cost of circuit board assembly, the lead time was reduced from 20 days to 5 days; the additional half hour of labor time reflects the dedication of a production line to the product. A similar analysis was performed for notebook assembly.

The three demand stages represent the delivery of product to the company’s retailers. The maximum service time for each of the demand stages equals zero. That is, they must provide immediate service to external customers. In the case of U.S. demand, the product can either be shipped by ground transportation at a cost of \$12 and a transportation time of five days or it can be shipped by air at a cost of \$20 with a one-day transportation time. Export demand can be satisfied in a similar manner, albeit with different costs and transportation times.

The current notebook is an improvement of an existing version. Therefore, the company used the previous product’s sales and market forecasts when determining the demand requirements for the supply chain. At each demand stage, the demand bound was estimated as

$$D_j(\tau) = \tau\mu + k\sigma\sqrt{\tau},$$

where μ and σ refer to the stage’s mean and standard deviation of demand and the constant k was chosen to equal 1.645. The daily demand parameters were estimated as shown in Table 2. The supply chain group felt that this demand bound captured the appropriate level of demand that they wanted to configure their system to meet using safety stock.

5.3. Different Solution Approaches

Table 3 summarizes the results from the four solution approaches evaluated in the project.

5.3.1. Minimum UMC Heuristic. The minimum UMC heuristic consists of choosing the lowest cost option at each stage and then optimizing the safety stock levels across the supply chain; the expected pipeline stock cost and COGS are the same as for the current policy, but the safety stock cost is lower. The optimal safety stock policy is to position several decoupling safety stocks across the supply chain, as shown in Figure 2, where each triangle denotes that the stage is holding safety stock.

Table 2 Demand Parameters for Different Markets

Demand stage	Mean	Sigma
U.S. demand—gray	200	120
Export demand—gray	75	50
U.S. demand—blue	125	80

Table 3 Results of Four Solution Approaches

	Current policy	Minimum UMC heuristic	Minimum lead-time heuristic	Supply chain configuration algorithm
Cost components (\$M)				
COGS	173.8	173.8	187.3	177.5
Pipeline stock cost	16.2	16.2	4.9	11.0
Safety stock cost	2.8	2.4	1.3	1.8
Total SC cost	192.7	192.3	193.5	190.4
Supply chain metrics				
Inventory investment	41.3	41.3	13.9	28.5
UMC (\$/unit)	1,737.56	1,737.56	1,872.93	1,775.43
Longest path (days)	91	91	35	68

All demand stages hold safety stock because they must quote a service time of zero. However, there is no safety stock at notebook or cover assembly; rather, the subassemblies that supply notebook assembly hold sufficient safety stock so that they can quote a zero service time for notebook assembly. For the circuit board, this translates into quoting a four-week service time for parts (which requires an inventory of eight-week parts).

The initial investment in safety stock and pipeline stock to create the supply chain equals \$41.4 million; the holding cost for safety stock and pipeline stock reflects the company’s 45% carrying cost. The expected demand over the course of one year is 100,000 units. Because a completed unit costs either \$1,737 or \$1,740, depending on the customer region, COGS dominates the total supply chain configuration cost.

5.3.2. Minimum Lead-Time Heuristic. The minimum lead-time heuristic chooses the single option at each stage with the shortest lead time. With the exception of electronic parts, the optimal placement of safety stocks is identical to the minimum UMC policy, but the actual stock levels differ due to the supply chains having different lead times. Because all the electronic parts are held on consignment, there is no need for a safety stock of eight-week parts, and it is optimal to hold inventory only at the subassembly and the finished goods stages.

Safety stock and pipeline stock costs are dramatically reduced due to the reduced lead times across the network. However, this comes with a 7.8% increase to the product’s UMC. The initial investment in safety stock and pipeline stock to create the supply chain equals \$13.8 million. The minimum lead-time heuristic results in a supply chain configuration cost that exceeds the minimum UMC heuristic by \$1.2 million.

5.3.3. Supply Chain Configuration Optimization.

In Table 4, we list the options selected by the algorithm presented in §4.2. In this configuration, we hold all electronic components on consignment, reduce the metal housing’s lead time to 30 days, and ship the finished goods by air. The optimal safety stock policy is represented in Figure 3.

The optimal policy holds a decoupling inventory at notebook assembly; this location allows demand pooling across the three end items. By placing electronics components on consignment and choosing the metal housing with the shortest lead time, the optimal solution is one where the upstream assemblies are

Figure 2 Optimal Safety Stock Placement for the Minimum UMC Heuristic

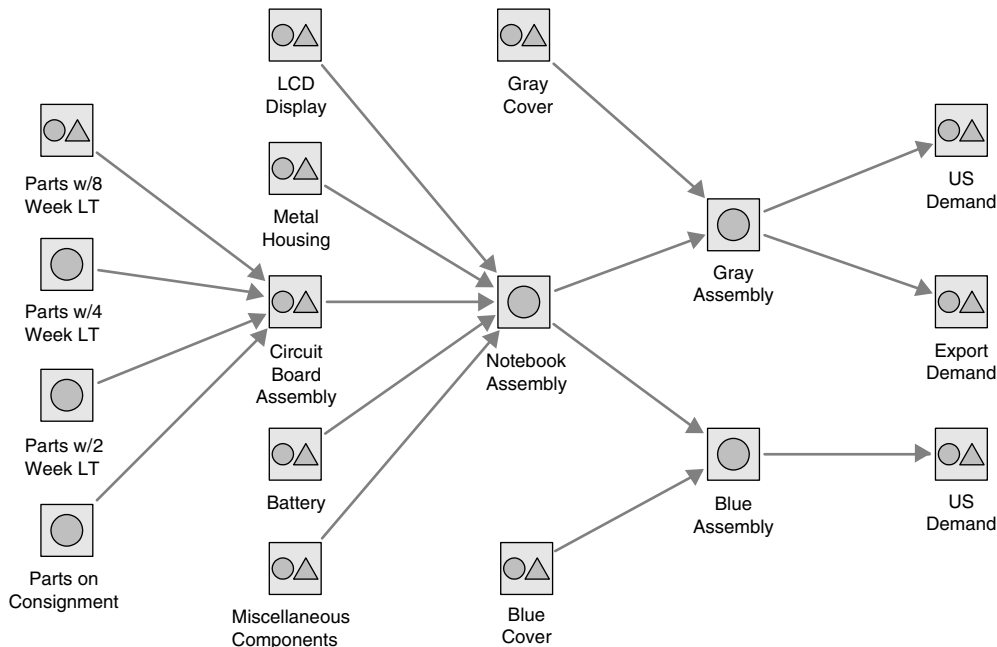


Table 4 Options Selected Using Optimization Algorithm

Component/process description	Option	Lead time	Cost added (\$)
Parts w/ eight-week lead time	4	0	136.59
Parts w/ four-week lead time	3	0	205.03
Parts w/ two-week lead time	2	0	156.93
Parts on consignment	1	0	200.00
Circuit board assembly	1	20	120.00
LCD display	1	60	300.00
Miscellaneous components	1	30	200.00
Metal housing	2	30	240.00
Battery	1	60	40.00
Notebook assembly	1	5	120.00
Gray cover	1	40	5.00
Blue cover	1	40	5.00
Gray assembly	1	1	30.00
Blue assembly	1	1	30.00
U.S. demand—gray	2	1	20.00
Export demand—gray	2	2	30.00
U.S. demand—blue	2	1	20.00

“balanced.” That is, each subassembly is configured to quote a service time of 30 days to the notebook assembly.

The initial investment in safety stock and pipeline stock equals \$28.5 million. This configuration increases the UMC by 2.2% over the minimum UMC heuristic but decreases the total supply chain cost by \$2.0 million, or \$20 per unit.

To put this cost savings in perspective, compare this to the savings from the current policy. The current policy chooses the cheapest option at each stage and holds inventory at all stages in the supply chain (as shown in Figure 1). The minimum UMC

heuristic optimizes the safety stock placement and saves \$0.4 million over the current policy. By optimizing the supply chain’s configuration, we save \$2.2 million, more than five times as much as optimizing just the safety stock placement.

Finally, increasing the UMC by \$37 is a nontrivial increase that would not be authorized without the kind of analysis given here. It is unlikely that the design team would ever discover this configuration by itself.

The optimal solution does not include some choices that one might have considered obvious. For example, one might be tempted to pay an additional \$0.50 to reduce the lead time for covers from 40 to 15 days. Similarly, \$5 to reduce battery lead time from 60 to 20 days might seem attractive. Yet in neither case does the benefit from the reduced lead time offset the increase in unit cost.

5.4. The Role of Holding Cost

An effective way for managers to strike a balance between COGS and inventory cost is through the choice of the holding cost rate. The holding cost rate can reflect not only the cost of capital and storage-related costs but also how much risk the company associates with making a large investment in safety and pipeline stock. This is particularly relevant for products like notebook computers that have short life cycles and high costs of obsolescence.

Figures 4 and 5 display inventory investment and supply chain configuration cost as functions of the holding cost rate for each of the three solutions

Figure 3 Optimal Safety Stock Placement for Optimization Algorithm

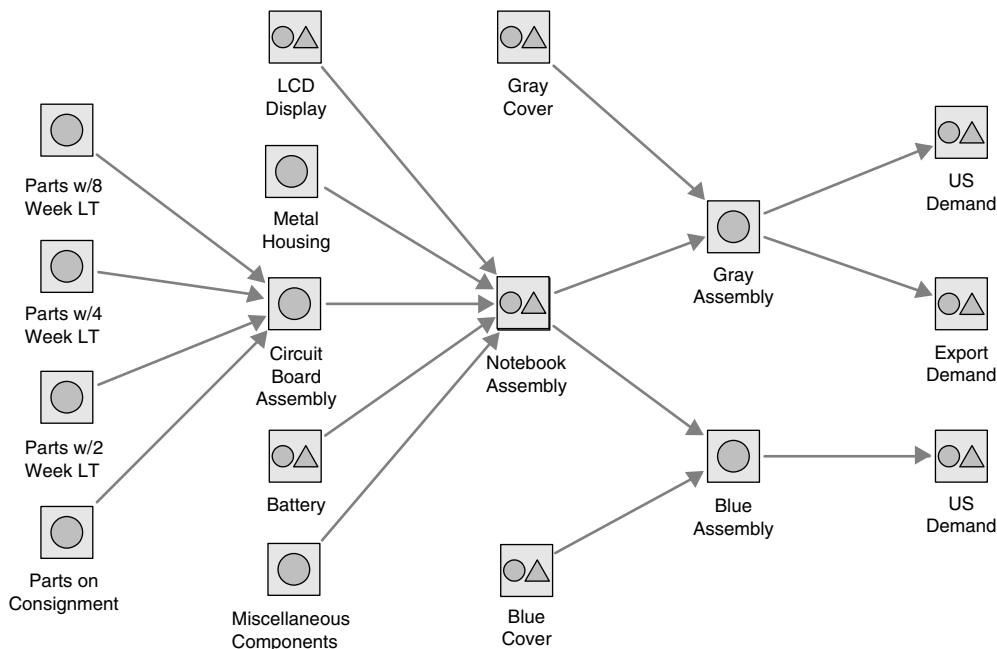



Figure 4 Inventory Investment as a Function of the Holding Cost Rate



approaches. In these two figures, we see how the optimal supply chain configuration moves from the minimum UMC solution to the minimum lead-time solution, as a function of the holding cost rate.

As the holding cost rate increases, the supply chain configuration algorithm chooses more higher-cost options, as seen in Table 5. When the holding cost rate is low, COGS dominates the total configuration cost. Therefore, the minimum UMC heuristic produces a solution that is very close to the optimal solution. But as the holding cost rate increases, the supply chain configuration algorithm creates a supply chain that comes closer to that for the minimum lead-time heuristic.

5.5. Insights Drawn From the Case Study

From this case study, we make five observations regarding supply chain configuration, and examine these observations in more detail in the following section.

OBSERVATION 1. In the optimal supply chain configuration, downstream stages are more likely to use high-cost options and upstream stages are more likely to use low-cost options and hold safety stock.

In the following observations, we make statements of the performance of the optimal supply chain configuration relative to the performance of the minimum UMC configuration. In §6, we also examine how

Table 5 Optimal Supply Chain Configuration Under Different Holding Cost Rates

	15%	30%	45%	60%	75%	90%
Parts w/ eight-week lead time	1	3	4	4	4	4
Parts w/ four-week lead time	1	2	3	3	3	3
Parts w/ two-week lead time	1	1	2	2	2	2
Parts on consignment	1	1	1	1	1	1
Circuit board assembly	1	1	1	1	2	2
LCD display	1	1	1	1	2	2
Miscellaneous components	1	1	1	1	1	1
Metal housing	1	1	2	2	2	2
Battery	1	1	1	1	1	2
Notebook assembly	1	1	1	2	2	2
Gray cover	1	1	1	1	1	1
Blue cover	1	1	1	1	1	1
Gray assembly	1	1	1	1	1	1
Blue assembly	1	1	1	1	1	1
U.S. demand—gray	1	2	2	2	2	2
Export demand—gray	2	2	2	2	2	2
U.S. demand—blue	1	2	2	2	2	2

these statements adapt in comparison to the minimum lead-time configuration.

OBSERVATION 2. The benefits of supply chain configuration increase as the importance of inventory costs increases relative to the total supply chain costs.

OBSERVATION 3. The benefits of supply chain configuration increase as the relative demand variability increases.

OBSERVATION 4. The benefits of supply chain configuration increase with longer lead times at downstream stages.

OBSERVATION 5. More echelons increase the benefits of supply chain configuration.

6. Serial-Line Design of Experiments

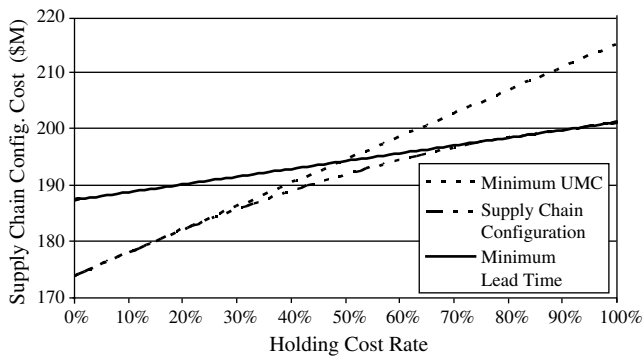
The purpose of the computational analysis in this section is to provide more evidence to support or refute the five observations made in the previous section. We first describe the test problem and then report results for each observation.

6.1. Scenario Definition

We consider a serial-line supply chain where there are low-cost and high-cost options at each stage. For each scenario, the selection of the low-cost option at each stage results in the total per unit cost of the product of \$100 and a total production time of 100 days.

We define a scenario by the number of stages in the serial line and a characterization of the high-cost option available at each stage. At each stage, the high-cost option is a fixed percentage more expensive and a fixed percentage faster than the low-cost option. Thus, scenario (8, 3, 30) denotes an eight-stage serial line where the high-cost option at each stage is 3% more expensive than the cost of the stage’s low-cost option

Figure 5 Supply Chain Configuration Cost as a Function of the Holding Cost Rate



and 30% faster than the production time of the stage's low-cost option.

For each scenario, we evaluate and solve 810 supply chain configuration problems that correspond to the permutations of three cost-accrual profiles, three time-accrual profiles, three mean demands, three standard deviations of demand, and 10 holding cost rates.

The three profiles for cost and time accrual are presented in Figure 6. We define the cumulative position x of stage i in an N -stage serial supply chain to be $x = i/N$; i.e., stage 1 is the raw material stage with $x = 1/N$, and stage N is the finished goods stage with $x = 1$. The profile $f(x)$ denotes the cumulative cost, or time, for all the low-cost options up to and including that stage. For both cost and time, we consider three profiles: $f(x) = x^{0.25}$, x , and x^2 . For example, if the cost profile were $f(x) = x^{0.25}$, then the cost of stage i 's low-cost option is

$$\left(\frac{i}{N}\right)^{0.25} - \left(\frac{i-1}{N}\right)^{0.25};$$

for the numerical results that follow, the profile calculations are rescaled to \$100 or 100 days and rounded to two significant digits.

These profiles allow us to capture the breadth of supply chain structures that exist in reality. For example, traditional consumer-electronics-manufacturing supply chains consist of high raw material cost and long component lead times (where both cost and time follow $x^{0.25}$), whereas an original equipment manufacturer (OEM) relying on outsourced manufacturing may see a similar cost profile but a time profile much more like x^2 , due to the fact that its supply chain responsibilities are primarily distribution based.

We consider high, medium, and low values for both the mean and standard deviation of demand. For both

the mean and standard deviation the respective values are 100, 50, and 10. Therefore, permutations will have coefficients of variation ranging from 0.1 to 10.

We consider 10 realizations of holding cost, varying from 10% to 100% in 10% increments. This range of values for holding cost captures different industries' assessment of the risk in their supply chains.

6.2. Hypothesis Evaluation

Unless otherwise noted, the results in this section consider an eight-stage serial supply chain. Figure 7 is a contour map representing the percentage difference between the optimal supply chain configuration (SCC) and the optimized minimum UMC heuristic.

To gain deeper insight into the problem space, we will focus our attention on the (8, 3, 30) scenario. For the 810 permutations associated with this scenario, on average the total cost of the minimum UMC heuristic exceeds the optimal SCC policy by 1.95%, and in 73% of the permutations the optimal SCC includes at least one high-cost option. The results from this scenario hold true across the range of scenarios displayed in Figure 7.

OBSERVATION 1. In the optimal supply chain configuration, downstream stages are more likely to use high-cost options and upstream stages are more likely to use low-cost options and hold safety stock.

Choosing a high-cost option at a downstream stage increases COGS but has only a local impact on pipeline stock cost. In contrast, choosing a high-cost option at an upstream stage increases not only COGS but also the safety and pipeline stock costs of all of its downstream stages. Second, holding safety stock at upstream stages is a relatively inexpensive way to

Figure 6 Cost and Time Profiles for Low-Cost Options




Figure 7 Contour Map for Scenarios with Eight-Stage Serial Supply Chain When the Percentage Changes for High-Cost Options' Cost and Time Vary From 1% to 100%

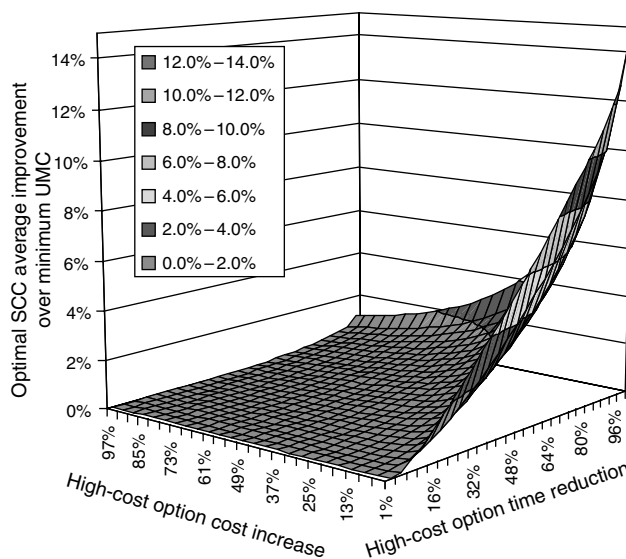
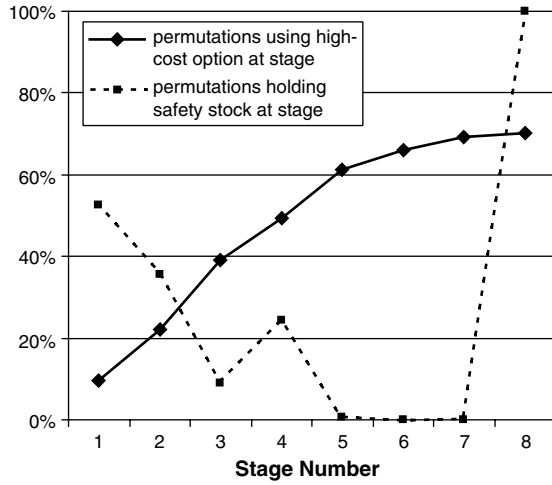


Figure 8 Percentage of Time Each Stage Holds Safety Stock and Selects the High-Cost Option (8, 3, 30 Scenario)



buffer long lead times at upstream stages. A safety stock acts to decouple the stage from the rest of the supply chain, making the effective lead time to the rest of the supply chain zero.

For the (8, 3, 30) scenario, Figure 8 plots the percentage of permutations (810 in total) for which each stage chooses the high-cost option and the percentage of permutations for which each stage holds a decoupling safety stock.

For this scenario, stages 6 and 7 never hold safety stock and stage 5 does so less than 1% of the time; by assumption, the finished goods inventory (FGI) stage, stage 8, provides immediate service to end item customers so it always holds safety stock. Downstream stages are much more likely than upstream stages to choose the high-cost option, while the upstream stages are more likely to hold inventory.

OBSERVATION 2. The benefits of supply chain configuration increase as the importance of inventory costs increases relative to the total supply chain costs.

When inventory costs are a small component of the total supply chain cost, then COGS dominates the

expression and the default solution to choose the lowest cost options will perform very well. As the inventory cost component of the total supply chain cost increases, there is an increasing likelihood that the inventory benefits from a high-cost option will outweigh the increase to COGS. When inventory costs are dominant relative to COGS, then the minimum lead-time heuristic will perform well relative to the optimal SCC.

In Table 6 we segment the (8, 3, 30) scenario by the holding cost rate and then compare the performance of SCC to that of the minimum UMC and minimum lead-time heuristics. When the holding cost rate is low, COGS dominates and the minimum UMC heuristic produces a solution that is very close to the optimal solution. The minimum lead-time heuristic performs best when the holding cost rate is high, and inventory costs are a more significant contributor to the total supply chain costs. Nevertheless, the optimal solution from the supply chain configuration algorithm differs from the heuristic solutions most of the time, especially for the mid ranges for the holding cost rates.

OBSERVATION 3. The benefits of supply chain configuration increase as the relative demand variability increases.

Greater demand variability results in a higher safety stock cost while leaving COGS and pipeline stock cost unchanged; thus, there is a greater likelihood that inventory savings from a high-cost option can offset the increase in COGS.

In Table 7 we segment the (8, 3, 30) scenario by the mean and standard deviation of demand. As the mean demand increases or the standard deviation of demand decreases, the benefit from optimizing the SCC decreases relative to the minimum UMC because the impact of COGS increases relative to the impact of safety stock. Conversely, relative to the minimum lead-time heuristic, as the mean demand decreases, or the standard deviation of demand increases, the benefit from optimizing the SCC decreases because the

Table 6 Scenario (8, 3, 30) Segmented by Holding Cost Rate

Holding cost rate (%)	Average amount min. UMC exceeds opt. SCC (%)	Percentage of permutations opt. SCC equals min. UMC	Average amount min. lead time exceeds opt. SCC (%)	Percentage of permutations opt. SCC equals min. lead time
10	0.1	75.3	2.2	0.0
20	0.4	58.0	1.7	0.0
30	0.7	28.4	1.4	1.2
40	1.2	27.2	1.2	3.7
50	1.7	21.0	1.0	3.7
60	2.2	17.3	0.9	6.2
70	2.6	16.0	0.8	7.4
80	3.1	14.8	0.7	11.1
90	3.6	7.4	0.6	13.6
100	4.0	2.5	0.6	17.3

Table 7a Average Percentage Amount Minimum UMC Exceeds Optimal SCC for (8, 3, 30) Segmented by Demand

	Low sigma (%)	Med sigma (%)	High sigma (%)
High mean	1.4	1.5	1.6
Med mean	1.4	1.6	1.9
Low mean	1.6	2.7	3.9

Table 7b Percentage of Permutations Where Optimal SCC Equals Minimum UMC for (8, 3, 30) Segmented by Demand

	Low sigma (%)	Med sigma (%)	High sigma (%)
High mean	34.4	33.3	30.0
Med mean	34.4	30.0	26.7
Low mean	30.0	13.3	8.9

Table 7c Average Percentage Amount Minimum Lead Time Exceeds Optimal SCC for (8, 3, 30) Segmented by Demand

	Low sigma (%)	Med sigma (%)	High sigma (%)
High mean	1.4	1.3	1.2
Med mean	1.4	1.2	1.0
Low mean	1.2	0.7	0.5

Table 7d Percentage of Permutations Where Optimal SCC Equals Minimum Lead Time for (8, 3, 30) Segmented by Demand

	Low sigma (%)	Med sigma (%)	High sigma (%)
High mean	1.1	2.2	3.3
Med mean	1.1	3.3	7.8
Low mean	3.3	14.4	21.1

impact of COGS decreases relative to the impact of safety stock.

OBSERVATION 4. The benefits of supply chain configuration increase with longer lead times at downstream stages.

Downstream stages have higher holding costs than upstream stages. As such, a lead-time reduction results in greater holding cost savings in pipeline and safety stock at a downstream stage relative to an upstream stage. In Table 8 we segment the (8, 3, 30) scenario by the cost-accrual and time-accrual profiles.

These tests confirm our hypothesis as the largest improvements from the SCC occur for the cases when the downstream lead times are greatest (T^2 time accrual function). In particular, the benefit from choosing the optimal SCC, relative to the minimum UMC heuristic, increases as more cost accrues upstream in the supply chain. As more cost accrues, the incremental cost of choosing the higher-cost options at downstream stages decreases, thereby making the high-cost option at a downstream stage relatively cheaper than in other configurations where more cost accrues downstream. This adds to the

Table 8a Average Percentage Amount Minimum UMC Exceeds Optimal SCC for (8, 3, 30) Segmented by Cost and Time Accrual

	$T^{0.25}$ (%)	T (%)	T^2 (%)
C^2	0.2	0.9	1.7
C	0.4	1.7	2.8
$C^{0.25}$	1.6	3.7	4.7

Table 8b Percentage of Permutations Where Optimal SCC Equals Minimum UMC for (8, 3, 30) Segmented by Cost and Time Accrual

	$T^{0.25}$ (%)	T (%)	T^2 (%)
C^2	73.3	37.8	17.8
C	68.9	17.8	7.8
$C^{0.25}$	17.8	0.0	0.0

Table 8c Average Percentage Amount Minimum Lead Time Exceeds Optimal SCC for (8,3,30) Segmented by Cost and Time Accrual

	$T^{0.25}$ (%)	T (%)	T^2 (%)
C^2	1.7	0.8	0.7
C	1.0	0.7	1.0
$C^{0.25}$	0.6	1.5	1.9

Table 8d Percentage of Permutations Where Optimal SCC Equals Minimum Lead Time for (8, 3, 30) Segmented by Cost and Time Accrual

	$T^{0.25}$ (%)	T (%)	T^2 (%)
C^2	5.6	0.0	0.0
C	15.6	0.0	0.0
$C^{0.25}$	36.7	0.0	0.0

attractiveness of choosing the higher-cost options at downstream stages.

The combination of these effects causes the 90 permutations in $[C^{0.25}, T^2]$ to have the greatest improvement because the high-cost options for downstream stages provide the largest “bang for the buck”; these options are cheap on a relative basis, yet yield a large reduction in lead time, which cuts most expensive inventory. All of the permutations that fall in this category use some high-cost options.

The comparison of the SCC to the minimum lead-time heuristic is also revealing. We see here that the minimum lead-time heuristic is never optimal when the time accrual function is T or T^2 , and the minimum lead-time heuristic seems to perform best when the cost and time profiles are similar. When the cost and time profiles are asymmetric, there will be some high-cost options in the minimum lead-time configuration that give very little benefit yet cost a significant amount.

OBSERVATION 5. More echelons increase the benefits of supply chain configuration.

More echelons provide a greater opportunity to offset an increase in COGS with a decrease in inventory costs because there are more configuration options from which to choose. In Table 9 we investigate the role that the number of echelons has in the benefits from optimizing the supply chain configuration.

Table 9 presents 16 different scenarios, each with 810 tests problems, corresponding to four different echelon structures. For each chain with fewer than eight echelons, we specify the times and costs for each stage so that the chain reflects the cumulative parameters of the eight-stage supply chain. For example, stage 1 in the two-stage supply chain has the cumulative cost and time for stages 1 to 4 in the eight-stage chain. In Table 9a we demonstrate that, for the

Table 9a Average Percentage Amount Minimum UMC Exceeds Optimal SCC for 16 Scenarios

Number of echelons	Option 2%		Option 3%		Option 4%		Option 5%	
	20%	30%	30%	40%	40%	50%	50%	
8	1.3%	2.0%	2.7%	3.6%				
4	1.2%	1.9%	2.7%	3.5%				
2	1.2%	1.8%	2.5%	3.3%				
1	0.9%	1.4%	2.0%	2.7%				

Table 9b Percentage of Permutations Where Optimal SCC Equals Minimum UMC

Number of echelons	Option 2%		Option 3%		Option 4%		Option 5%	
	20%	30%	30%	40%	40%	50%	50%	
8	27.2%	26.8%	26.8%	26.7%				
4	25.3%	24.2%	23.8%	23.8%				
2	24.9%	24.8%	24.0%	22.8%				
1	35.6%	35.6%	35.6%	35.6%				

Table 9c Average Percentage Amount Minimum Lead Time Exceeds Optimal SCC for 16 Scenarios

Number of echelons	Option 2%		Option 3%		Option 4%		Option 5%	
	20%	30%	30%	40%	40%	50%	50%	
8	0.8%	1.1%	1.4%	1.8%				
4	0.7%	1.1%	1.4%	1.7%				
2	0.6%	0.9%	1.2%	1.4%				
1	0.3%	0.5%	0.6%	0.7%				

Table 9d Percentage of Permutations Where Optimal SCC Equals Minimum Lead Time

Number of echelons	Option 2%		Option 3%		Option 4%		Option 5%	
	20%	30%	30%	40%	40%	50%	50%	
8	6.2%	6.4%	7.2%	8.0%				
4	9.1%	10.4%	11.5%	12.5%				
2	18.8%	20.1%	22.2%	23.5%				
1	64.4%	64.4%	64.4%	64.4%				

same option-cost structure, more echelons increase the benefits of optimizing the supply chain's configuration. It is also true that in the framework of this design of experiments, more echelons provide us with stages that have cheaper options relative to chains with fewer stages. This increases the likelihood that some stages will yield a benefit from using the available high-cost option.

7. Conclusion and Next Steps

In this paper, we introduce and develop a model for configuring new supply chains. We model the supply chain as a network whose nodes represent functional requirements in the supply chain. For each node, multiple options can exist to satisfy the functional requirement. The optimal supply chain configuration minimizes the sum of three relevant costs: COGS, safety stock cost, and pipeline stock cost.

As a form of validation, we describe a diagnostic exercise where the model was used to determine and evaluate the best sourcing strategy for a notebook computer supply chain. This case study demonstrated that cost savings can be realized when inventory cost and COGS are jointly optimized. Prior applications of this approach are described in Graves and Willems (2000b) and Wala (1999). We state some general observations based on these examples and confirm the observations by means of a computational experiment.

This research raises several relevant questions for further consideration; we conclude with five issues that we think are most worthy of additional work. First, *time-to-market costs* should be incorporated. This can be accomplished by augmenting the two-state formulation with a third state variable. The additional state variable is the maximum replenishment time (corresponding to the longest path) in the network. Once the state variable has been added, the time-to-market cost can be included in the criterion function. Second, *more general network structures* should be considered. Spanning trees can capture commonality for a few critical parts or processes, but there is a definite need to allow more general component and process commonality. Third, many firms wish to *limit the number of different vendors* that the model can choose. We suspect that the most fruitful approach will be to preprocess the available options at a stage, pruning some options that would violate the constraints on additional vendors. If this were done as a preprocessing step, then the original solution procedure will remain valid.

Fourth, we assume that a single option is chosen at each stage. In reality, a firm will often decide to have *dual or multiple sources*. For instance, a firm might opt to have both a cheap, long lead-time supplier

and a more expensive, short lead-time supplier. The long lead-time supplier would be the source for the more predictable portion of demand, whereas the less predictable demand would be served by the short lead-time supplier. We have seen a similar strategy applied in choosing transportation modes (Threatte and Graves 2002) for product shipments from Asia to North America, where a mixed strategy of using both air and ship provides an effective hedge against demand uncertainty.

We can think of two possible ways to incorporate multiple sources in the configuration problem, neither of which is very satisfactory. First, one might decide a priori to have, say, two sources and to split the demand according to fixed percentages, say 60% and 40%, between the two sources. Then, one can replace this stage in the network model with two stages, where one stage gets 60% of the demand and the other gets 40% of the demand. We would need to identify options for each of these two stages and adapt the algorithm to ensure that the same supply source is not chosen for each option. This approach would provide accurate estimates of the pipeline inventory costs and COGS; however, the current model would not provide an accurate estimate of the safety stock as it would think that the two stages provide different components in fixed proportions of 60% and 40%. A variant of this approach would be to split the demand by variability, i.e., a predictable, base level of demand and an uncertain component, and then source each separately.

A second approach is to create options that represent the choice of dual or multiple sources for a stage. For instance, suppose that there are two options, A and B; Then, we could create a third option representing dual sourcing from both A and B. Assumptions would be needed to decide the cost and lead-time inputs for this third option, so that the model would be able to determine the configuration costs; again, the most challenging aspect would seem to be how to get a good estimate for the safety stock requirements.

The fifth consideration is to develop the supply chain configuration problem for more conventional assumptions on the behavior of the inventory system. In particular, we assume a guaranteed service model; that is, each stage provides 100% service for its quoted service time. To determine the safety stock required for this guarantee, we assume that demand is bounded. In contrast, much of the multi-echelon inventory literature assumes a stochastic service model in which the service time between stages can vary depending upon material availability at the upstream stage. For the stochastic service model, the determination of the safety stock typically requires knowledge of the demand distribution. Graves and

Willems (2003) discuss and compare these two models in the context of the safety stock placement problem.

We can suggest a few approaches to the supply chain configuration problem with the stochastic service model. One approach would be an iterative algorithm that would iterate between an optimization that selects the options and an optimization for the safety stock placement. The first optimization could be an integer linear program that chooses the options so as to minimize COGS and the pipeline holding cost, given the safety stocks. The second optimization would be a multiechelon inventory model that minimizes the cost of the safety stock for the set of selected options. Open questions are how to ensure that such a procedure would converge and whether it would converge to a global optimum. A second approach might be to develop an analytic expression for the safety stock as a function of the replenishment time at a stage. This expression would need to account for the stochastic nature of the replenishment time for the stage and its dependence on the choice of options at upstream stages. As such, this is likely to require an approximation. But then, if one can model the safety stock as a function of the replenishment time, one could conceivably use the dynamic program given in this paper, or some variant of this.

Acknowledgments

This research was supported in part by the Massachusetts Institute of Technology (MIT) Leaders for Manufacturing Program, a partnership between MIT and major manufacturing firms, and by the Singapore-MIT Alliance, an engineering education and research collaboration among the National University of Singapore, Nanyang Technological University, and MIT. The second author was supported by the Boston University School of Management Junior Faculty Research Program. The authors thank the editors and referees for their very helpful comments and advice on earlier versions of the paper.

References

- Ansari, S. L., J. E. Bell. 1997. *Target Costing: The Next Frontier in Strategic Cost Management/The CAM-I Target Cost Core Group*. Irwin, Chicago, IL.
- Arntzen, B. C., G. G. Brown, T. P. Harrison, L. L. Trafton. 1995. Global supply chain management at Digital Equipment Corporation. *Interfaces* 25(1) 69–93.
- Ettl, M., G. E. Feigin, G. Y. Lin, D. D. Yao. 2000. A supply network model with base-stock control and service requirements. *Oper. Res.* 48(2) 216–232.
- Geoffrion, A. M., G. W. Graves. 1974. Multicommodity distribution system design by Benders decomposition. *Management Sci.* 20(5) 822–844.
- Geoffrion, A. M., R. F. Powers. 1995. Twenty years of strategic distribution system design: An evolutionary perspective. *Interfaces* 25(5) 105–127.
- Graves, S. C., S. P. Willems. 2000a. Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1) 68–83.

- Graves, S. C., S. P. Willems. 2000b. Optimizing the supply-chain configuration for new products. *Proc. 2000 MSOM Conf.*, Ann Arbor, MI.
- Graves, S. C., S. P. Willems. 2003. Supply chain design: Safety stock placement and supply chain configuration. A. G. de Kok, S. C. Graves, eds. *Handbooks in OR & MS*, Vol. 11. Elsevier, Amsterdam, The Netherlands.
- Inderfurth, K. 1993. Valuation of leadtime reduction in multi-stage production systems. G. Fandel, T. Gullledge, A. Jones, eds. *Operations Research in Production Planning and Inventory Control*. Springer, Berlin, Germany, 413–427.
- Threatte, K., S. C. Graves. 2002. Tactical shipping and scheduling at Polaroid with dual lead-times. *Proc. 2002 SMA Conf.*, Singapore.
- Wala, T. 1999. Build-to-order/direct ship model for professional digital cameras. Unpublished S. M. dissertation, Leaders for Manufacturing Program, Massachusetts Institute of Technology, Cambridge, MA.
- Willems, S. P. 1999. Two papers in supply-chain design: Supply-chain configuration and part selection in multigeneration products. Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA.