Interfaces
Vol. 37, No. 5, September–October 2007, pp. 420–435
ISSN 0092-2102 | EISSN 1526-551X | 07 | 3705 | 0420

informs®

DOI 10.1287/inte.1070.0298

A Periodic-Review Modeling Approach for Guaranteed Service Supply Chains

John M. Bossert
Optiant, Incorporated, Burlington, Massachusetts 01803, john.bossert@optiant.com

Sean P. Willems
School of Management, Boston University, Boston, Massachusetts 02215, willems@bu.edu

We extend the guaranteed service, supply chain modeling framework to allow for an arbitrary, integer review period or ordering frequency at each stage. We define a notation for the cyclic inventory dynamics that review periods introduce and generalize inventory-balance equations to accommodate three different periodic-review operating policies—constant base stock, constant safety stock, and adaptive base stock. As a form of validation, we apply the model to the Celanese acetic acid supply chain and show that inventory metrics of the new model differ by more than 30 percent from those derived through the simpler modeling approach of aggregating a review period into lead time.

Key words: multiechelon; inventory system; safety stock; optimization; dynamic programming application; review periods.

Leading-edge companies have integrated lean manufacturing and Six Sigma processes deeply into their organizations to reduce the total length and cost of their supply chains while maintaining or increasing service to their customers. Against this backdrop, firms are also outsourcing significant portions of their operations and integrating their processes more tightly into the delivery networks of their customers. An integrated and long-lasting supply chain improvement process increases inventory turns, increases return on assets, and decreases cash-to-cash cycle times. These factors improve corporate performance. However, continued advances in lean initiatives and the increasing complexity of supply chains pose modeling challenges. In particular, these improvements increase the importance of correctly modeling operating times and policies across the supply chain; they raise the bar that a useful model must cross.

Periodic-review models make up a rich research stream within the field of inventory management; for a broad survey of such models, we refer the reader to Federgruen (1993). However, the complexity that review periods introduce seems to have limited results for chains with stage-dependent review periods to specialized systems. Papers illustrative of this work include Graves (1996) and van Houtum et al. (2003). Graves (1996) develops a computationally intensive, exact evaluation approach of inventory levels in a distribution system facing demands with independent increments. Van Houtum et al. (2003) prove that a base-stock policy is optimal for a serial-line network with nested review periods.

Developing a model to optimize inventory levels and locations across a supply chain in the presence of review periods entails trading off exactness and tractability. Our approach is similar in spirit to that of Lee and Billington (1993), who develop a decentralized model to set inventory levels across a multiechelon supply chain subject to demand and supply variability. To make their model tractable, the authors richly characterize a single-stage inventory model; however, they assume that the demand process each stage faces is an allocation of the end-item demands that the stage satisfies.

In this paper, we extend the discrete-time, supply chain modeling framework that Simpson (1958) originally described to allow for an arbitrary, integral review period at each stage of a chain as well as different inventory policies. Simpson (1958) defined the guaranteed service modeling (GSM) framework for a serial-line and distribution network. Graves and Willems (2000), which we hereafter refer to as GW, extended the framework to supply chains that are modeled as spanning trees; they formulated a deterministic, dynamic program to optimize the spanning tree models. Humair and Willems (2006) further generalized the network structure to so-called clusters of commonality. Application of this modeling approach at Hewlett-Packard was a 2003 Edelman Prize finalist (Billington et al. 2004). All previous GSM work assumes a single underlying review period that is common to all stages. We summarize the GSM framework in Appendix 1—*Reviewing the GSM Framework*.

Although single-stage models often readily accommodate review periods, stage interactions can greatly complicate multiechelon models. First, review periods complicate demand propagation. A stage that reviews periodically typically orders periodically and so generates intermittent demand. Intermittent or more general, nonstationary incoming demand, combined with periodic review, compounds the complication. Although nested review periods effectively negate intermittency and seem broadly appropriate, they do not always appear in practice, as the real-world example in the *Application at Celanese* section illustrates. Without nesting, one must account for not just review-period lengths but also for staggering. A stage that orders every two days must distinguish weekly demand originating on Mondays from weekly demand originating on Fridays. In addition, many different ordering policies exist. A stage might always order up to a fixed, precalculated base-stock target. This case seems most straightforward, and we consider it first. Alternatively, a stage might order to maintain a constant safety-stock level, choose a fully adaptive base-stock policy, or even smooth demand. Ordering behavior affects inventory dynamics at the stage in question and further complicates demand propagation. In the *Extension for General Review Periods* section, we describe the models and briefly generalize the dynamic GW program to accommodate review periods.

The *Application at Celanese* section demonstrates the importance of review periods by presenting the successful application of this model at Celanese, a $6 billion chemical company. Celanese and the chemicals industry in general encompass a host of review-period variations that are often too critical to ignore. Boats operating under fixed schedules transport many raw materials and finished goods. Customers are assigned specific days to order each week, and some distribution centers review at different frequencies. Sometimes requirements are simply transmitted monthly, and sometimes they are smoothed over the monthly review cycle. Finally, the capital intensity of the business makes cyclic schedules commonplace. Our modeling framework addresses each of these issues. Although we focus on the application at Celanese, we have integrated our review-periods research into the Optiant PowerChain software application. More than a dozen Fortune 500 companies, including Black and Decker, Boston Scientific, Hewlett-Packard, Honeywell, Intel, and Procter & Gamble, have applied it.

We offer some conclusions in the *Conclusion* section.

Extension for General Review Periods

Extending the GSM framework to include review periods involves two primary complications—characterizing internal demand streams and generalizing the inventory-balance equation. In a single-stage setting (Hadley and Whitin 1963), the time interval of interest is the order cycle that elapses between consecutive orders or consecutive replenishments. The multiechelon setting involves three cycles at each stage. The order cycle still exists and still equals the review period. However, this cycle operates in concert with two additional cycles. First, incoming demand might be intermittent or more generally cyclic, and the cycle length of inbound demand depends on downstream review periods. A third cycle governs inventory dynamics at the stage itself as well as outgoing demand transmission to its suppliers. We assume that the review period defines the frequency with which a stage places demands on its suppliers or, more generally, modifies its ordering behavior. In addition, the suppliers receive demand information only through these orders, although they know their customers' inventory policies.

The *Demand Propagation Under a Constant Base-Stock Target* and *Single-Stage Model* sections develop the notation and inventory-balance equation for a single stage that resets its inventory position to a constant base-stock target in the presence of arbitrary review periods. Using two examples, the *Example* section illustrates the inventory dynamics that the balance equation implies. The *Demand Bounds and Service-Level Targets* section connects the demand bounds to service-level targets, and the *Optimization* section generalizes the GW dynamic program to account for review periods. The *Adaptive Base-Stock and Constant Safety-Stock Policies* and *Demand Smoothing* sections extend the analysis to adaptive base-stock targets and a particular version of smoothing.

Demand Propagation Under a Constant Base-Stock Target

If a stage has a constant base-stock target, its ordering process under review periods remains simple; each review period, it resets its inventory position to the target by ordering the demand incurred since it last reviewed. We denote the length of the stage j review period by R_j and a corresponding offset by $\omega_j \in \{0, 1, 2, \ldots, R_j - 1\}$. That is, stage j places orders at times $\omega_j + n \cdot R_j$ for $n = 0, 1, 2, \ldots$. Offsets permit discrimination among stages that, for example, review weekly but on different days. Although the external demand processes remain stationary, stage-dependent review periods make the internal demands cyclic, and we denote the length of the demand cycle that stage j faces by λ_j^{in}. That is, for integers n and some fixed time t, the demands that stage j faces at times $t + n \cdot \lambda_j^{\text{in}}$ are independent and identically distributed. Because stage j might not order every period, the cycle length of the demand process that stage j generates might differ from that of the demand process that it faces. In particular, this outgoing cycle length is the least common multiple of the incoming cycle length and stage j's review period. We denote the length of stage j's outbound demand cycle by $\lambda_j^{\text{out}} = LCM(\lambda_j^{\text{in}}, R_j)$. Similarly, the inbound demand-cycle length at stage j is the least common multiple of the outbound demand-cycle lengths generated by the stages immediately downstream from stage j. That is, $\lambda_j^{\text{in}} = LCM(\{\lambda_k^{\text{out}} \mid k\colon (j,k) \in A\})$. We can calculate inbound and outbound demand-cycle lengths by first considering the stages facing external demand, then the stages upstream to only demand-facing stages, and so on.

Given the demand process that stage j faces, we can calculate the demand it places upstream at some time t by summing incoming demands over the relevant time window of R_j consecutive periods. That is, $d_j^{\text{out}}(t) = d_j(t - R_j, t)$ if $t = \omega_j + n \cdot R_j$ for some integer n, and zero otherwise. Given the cyclic demand processes of the stages immediately downstream from a stage j, we can characterize the demand that stage j faces in some period t of its demand cycle as $d_j(t) = \sum_{k \mid (j,k) \in A} \phi_{jk} \cdot d_k^{\text{out}}(t)$. We can propagate demand throughout the chain by again starting with stages facing external demand and proceeding upstream. Like GW, we assume that demand bounds $D(\cdot,\cdot)$ exist. However, because demand is now cyclic, each bound specifies a window of time rather than just a length. More specifically, we assume the existence of functions $D(\cdot,\cdot)$ such that $D_j(t,\tau) \geq \sum_{s=1}^{\tau} d_j(t+s)$ for $\tau > 0$ and $t \in \{1, 2, \ldots, \lambda_j^{\text{in}}\}$.

Single-Stage Model

We next develop the inventory-balance equation for a single stage; for clarity, we omit stage indices from this section. The net inventory-balance equation at time t generalizes to

$$I(t) = B - d(t - SI - T - x(t), t - S), \tag{1}$$

where B is the constant base-stock target, and $x(t) = (t - T - SI - \omega) \bmod R$. The equation differs from that of GW by only the $x(t)$ term, and if $R = 1$, the equation reduces to that of GW. The first argument to $d(\cdot,\cdot)$, $t - SI - T - x(t)$, corresponds to the last demand replenished by the stage's supplier by time t, and the correction term $x(t)$ reflects the additional inventory exposure because of review periods that are greater than one. Also, $x(\cdot)$ affects increments of the first argument of $d(\cdot,\cdot)$ in multiples of the review period R; the stage receives a replenishment only once every R periods. To derive the expression for $x(\cdot)$, note that the last replenishment order to arrive in inventory by time t replenished demands through the corresponding order time, say $\omega + n^* \cdot R$, where n^* is the largest integer such that $\omega + n^* \cdot R + T + SI \leq t$, assuming that a replenishment is available to serve demand in its period of arrival. The first argument to $d(\cdot,\cdot)$ equals the time $\omega + n^* \cdot R$, and $x(t)$ is the quantity that makes the above inequality defining n^* an equality. The latter argument of $d(\cdot,\cdot)$, $t - S$, corresponds to the last demand fulfilled by time t.

The distribution of the inventory level $I(t)$ is cyclic with cycle length λ^{out}. The length of the demand window, $d(t-SI-T-x(t), t-S)$, cycles from $T+SI-S$ to $T+SI-S+R-1$ as $x(t)$ cycles from 0 to $R-1$. The incoming demand stream has cycle length λ^{in}, and the demand window advances along this stream R units every review period. In turn, the window contents and the inventory level have the same distribution every $LCM(R, \lambda^{\text{in}}) = \lambda^{\text{out}}$ periods. We assess average inventory and other metrics over λ^{out} consecutive periods. Also, we generalize the definition of net-replenishment lead time (see Appendix 1) to $T+SI-S+R-1$, the maximum inventory-exposure length.

We set the base-stock target B to the smallest demand bound that guarantees a nonnegative inventory level $I(t)$. The inventory level cycles, so we set the target to the greatest bound on the demand term of the inventory-balance equation. Because $D(\cdot, \tau)$ increases with τ, we can restrict the search to times when inventory is exposed to a full net-replenishment lead time of $t+SI+T-S+R-1$ periods. In turn, the base stock should be set to the largest demand bound corresponding to times $\omega + n \cdot R + T + SI - 1$ for $n = 0, 1, 2, \ldots, \lambda^{\text{out}}/R - 1$, the times just before orders arrive. To see this result from another perspective, consider some time $t = \omega + n \cdot R$. The order placed at t will arrive in inventory at time $t + SI + T$, and the subsequent order will arrive at time $t + SI + T + R$. Consequently, the base stock needs to cover demand over $(t, t+SI+T-S+R-1]$, and the single base-stock target is the greatest demand bound corresponding to such an interval. Therefore, by either argument,

$$B = \max\left\{D(t, SI+T-S+R-1) \mid t = \omega + n \cdot R,\; n = 0, 1, 2, \ldots, \frac{\lambda^{\text{out}}}{R} - 1\right\}.$$

Example

We next illustrate the dynamics of the review-periods model using two examples. Consider first the two-stage chain of Figure 1.

Stage B supplies stage A, and the review periods are nested because the upstream review period of 4 is a multiple of the downstream review period of 2. External demand is normally distributed with mean and standard deviation of 10 units per period.

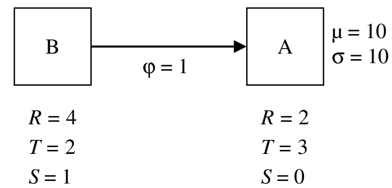

Figure 1: We illustrate a two-stage chain with nested review periods.

Stages A and B have processing times of 3 and 2, respectively, and quote service times of 0 and 1. The arc multiplier is 1, and neither stage has a positive offset ω. During every other period, stage A places an order on stage B that is normally distributed with mean 20 and standard deviation $\sqrt{10^2+10^2} = 14.14$. An order placed by A arrives in its inventory four periods after placement: one period-of-service time from B and three periods-of-processing time at Stage A. Therefore, given its review period of two periods, stage A is exposed alternately to four and five periods of demand. Also, $\lambda_A^{\text{out}} = \lambda_B^{\text{in}} = 2$. Demand is placed on stage B during every other period. However, because it orders every four periods, its demand is effectively stationary. Each order that stage B places has mean 40 and standard deviation $\sqrt{400} = 20$, and it arrives in stage B inventory two periods after placement, assuming the external supplier to stage B has a service time of 0.

Next, suppose we add to the chain a second demand stage C that reviews every three periods and receives external demand with mean and standard deviation of 20 units per period, as Figure 2 shows. The review periods are no longer nested because the stage B review period of 4 is not a multiple of 3.

Although stage C reviews every three periods, its inventory dynamics are qualitatively similar to those of stage A because both receive external demand. However, the stage C review period complicates the dynamics at stage B. Stage B still receives an order from A every other period; it now receives an order from C every third period, and $\lambda_B^{\text{in}} = LCM(2,3) = 6$. Stage B still orders every four periods. However, the demand stream it sees is no longer effectively stationary, and the stage B inventory dynamics operate on a $\lambda_B^{\text{out}} = LCM(6,4) = 12$-period cycle (Table 1).

Table 1 charts stage B activity over time and includes demands from A and C, as well as orders

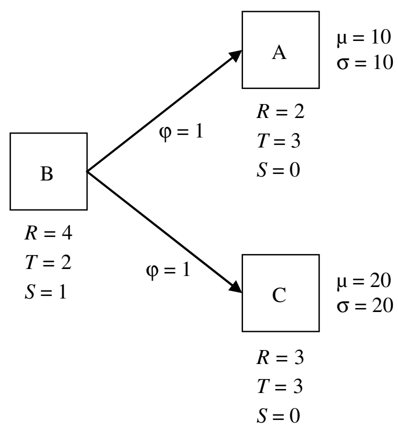

Figure 2: We show a three-stage supply chain with nonnested review periods.

placed by B every fourth period. Demands placed on stage B are fulfilled a service time $S_B = 1$ after receipt. In addition, stage B orders arrive in its inventory two periods after placement. This translates to periods 2, 6, and 10 of each 12-period cycle. For any period t, stage B inventory is exposed to demands received over the time window $(t - SI - T - x(t), t - S] = (t - x(t) - 2, t - 1]$, and the maximum exposure length of four periods occurs when $x(t) = 3$, at times 1, 5, and 9, the periods immediately before a replenishment arrival. At period 1, stage B is exposed to two demands from A and two from C, and at times 5 and 9, the exposure includes two demands from A and just one from C. Therefore, the stage B constant

Period	Stages placing demands	Fulfilling demand to stages	Stage B placing order	Stage B receiving order	
1		A, C			
2	A			Yes	
3	C	A			
4	A	C	Yes		
5		A			
6	A, C			Yes	
7		A, C			
8	A		Yes		
9	C	A			Stage B's largest exposure spans the last four periods of its 12-period cycle.
10	A	C		Yes	
11		A			
12	A, C		Yes		

Table 1: Table data illustrate stage B inventory dynamics.

base-stock target corresponds to the bound on two orders each from stages A and C.

Demand Bounds and Service-Level Targets

Although our treatment of review periods assumes no demand distribution, distributional assumptions are convenient, if not necessary, to actually calculate base-stock values. Indeed, in practice we consider demand bounds only implicitly and propagate a service target to each stage. If each review period is one, each stage faces stationary demand, and GW noted that assuming stage j receives normally distributed demand with mean μ and standard deviation σ implies a bounding function of the familiar form $D(\tau) = \tau \cdot \mu + z \cdot \sigma \cdot \sqrt{\tau}$, where τ is the net-replenishment lead time, and z is a safety factor. Under the guaranteed service assumption, we can readily translate each such bound into a Type I service level, the probability of fulfilling all demand of a given period on time.

Under review periods, we generalize such implied service-level targets from the Type I definition. A base stock that ensures a sufficiently low stock-out probability for each period of λ^{out} consecutive periods would have inconsistent meaning. For example, for a stage that receives stationary demand, 95 percent service under a review period of 4 will be greater than 95 percent service under a review period of 1; in the former case, the stock-out probabilities in the three days after an order arrives are likely to be near zero percent. Consequently, we calculate base-stock targets that achieve a weighted average of the Type I service levels defined over an inventory cycle of length λ^{out}, with one value taken per base time unit. Because we have no direct means of inversing this composite service level, we employ iterative searches to find the base-stock target that achieves the desired service level.

Optimization

We optimize service times in supply chains with review periods by applying the dynamic program of GW to a generalized version of problem (P) (see Appendix 1). More specifically, we calculate inventory levels through Equation (1) rather than the inventory balance of GW, and we generalize the first set of constraints to $S_j - SI_j \leq T_j + R_j - 1$ to reflect the generalized definition of net-replenishment lead time. Although the inventory level of GW is stationary, we must average the expected inventory levels implied by Equation (1) over the λ^{out} periods of an inventory cycle; to increase computational tractability, we make several approximations. We approximate the demand term of Equation (1), $d(t - SI - T - x(t), t - S)$, by assuming stationary demand. More specifically, we note the length of the delimited time window cycles between $SI + T - S$ and $SI + T - S + R - 1$, and we make the approximation

$$E\left[\frac{1}{\lambda^{\text{out}}} \cdot \sum_{t=1}^{\lambda^{\text{out}}} d\left(t - SI - T - x(t), t - S\right)\right] \approx \left(\frac{1}{\lambda^{\text{out}}} \cdot \sum_{t=1}^{\lambda^{\text{out}}} \mu_t\right) \cdot \left(SI + T - S + \frac{R-1}{2}\right) \cdot R,$$

where μ_t is the mean of time t demand.

In addition, we do not evaluate the base-stock term of Equation (1) at all service times exactly because each evaluation entails the relatively expensive iterative search that we note in the *Demand Bounds and Service-Level Targets* section. We make the approximation that the base stock is a function of the times SI and S through only the difference $SI - S$, evaluate the base stock exactly at several values of $SI - S$, and then apply interpolation to evaluate intermediate values.

More significantly, review periods introduce dependence between demand propagated to a stage and downstream service times; the dynamic program cannot readily address this circularity. By contrast, demand propagation under the GW model is independent of service times and can be performed before optimization. Minimizing inventory levels under the generalized version of problem (P) would require repropagating demand at each stage of the dynamic program; this would, in turn, involve determining the optimal cost-to-go solution and break the efficiency of dynamic programming. Consequently, we propagate a priori the demand stream for service times equal to zero and then use this stream throughout the dynamic programming algorithm.

We tested the above approach against exhaustive enumeration on 189 five-stage serial-line supply chains. The 189 chains represent the permutations of three cost-accrual profiles, three lead-time profiles, and 21 review-period profiles. In 131 of the 189 chains, the approximate solution generated by this section produces the same solution as the exact solution found by enumeration. The average error across all chains is 1.27 percent; for the 58 chains where the approximation did not produce the optimal safety-stock policy, the conditional average error was 4.14 percent (see Appendix 2 for details). These results and the real-world demand for efficient computation support the use of the approximation approach.

Adaptive Base-Stock and Constant Safety-Stock Policies

A constant base-stock policy implies ordering, at each review period, the demand incurred since the previous review period. Therefore, it increases clarity of analysis. We next derive the ordering behavior under a general adaptive base-stock policy, assuming a set of precalculated targets, one for each review period over an outgoing demand cycle of length λ^{out}. Such a policy might lower inventory levels or otherwise improve operations. By construction, a stage orders up to its base-stock target at each review period. Now, however, the targets vary over time. The effect on Equation (1) is that the base-stock term now requires a time index. We find it convenient to decompose the base-stock target into the expected demand that the target must cover and safety stock. Fixing the latter term defines a constant safety-stock policy, a special case of adaptive policies that we consider at the end of this section.

We express the target as

$$B(t) = \sum_{s=t+1}^{t+SI+T+R-1-S} \mu_s + SS(t).$$

Recall from the *Single-Stage Model* section that an order placed at review period t effectively covers demand over $(t, t+SI+T-S+R-1]$. The latter term of $B(t)$ is a potentially adaptive safety-stock value, and it corresponds to the expected inventory level at time $t + SI + T + R - 1$, immediately before arrival of the following order.

Because a stage still orders up to its base-stock target, we can calculate the quantity ordered at review period t, $O(t)$, as

$$O(t) = B(t) - IP(t), \tag{2}$$

where $IP(t)$ is the inventory position at time t, after demand arrival but prior to order placement. As such, $IP(t)$ is

$$IP(t) = B(t-R) - d(t-R, t). \tag{3}$$

Substitution yields

$$O(t) = d(t-R, t) + B(t) - B(t-R). \tag{4}$$

Qualitatively, the order consists of two parts—the demand observed since the previous review period as under a constant target, as well as a base-stock adjustment. The adjustment terms telescope over time, and their sum over a cycle of length λ^{out} is zero. In addition, the adjustment terms permit negative order quantities. A real operation would be unlikely to disassemble goods and ship them upstream; we assume that the distortion of a stage retaining a negative demand instead of returning it upstream is small enough to ignore safely.

Even without product disassemblies, the regular target adjustments of an adaptive base-stock policy can prove difficult to implement, particularly through an enterprise advanced planning and scheduling (APS) software application that requires specific inventory targets as input. A constant safety-stock policy seems easier to implement and more in concert with APS systems. Under a constant safety-stock policy, we eliminate the time index from the safety-stock portion of the base-stock target:

$$B(t) = \sum_{s=t+1}^{t+SI+T+R-1-S} \mu_s + SS.$$

Under a fully adaptive base-stock policy, we might be able to express $SS(t)$ in terms of demand-distribution parameters. In particular, under normal demands, setting $SS(t) = z \cdot \sqrt{\sum_{s=t+1}^{t+SI+T+R-1-S} \sigma_s^2}$ achieves the Type I service level corresponding to z in each review period of length R. Under constant safety stock, however, we again use a search procedure to find the constant value that makes the composite service level over an inventory cycle sufficiently high.

Demand Smoothing

The above models assume that, although a stage might incur a demand during each period, it receives a single replenishment order each review period. Therefore, upstream demand can easily become very spiky. An inventory manager might want to review only periodically, yet propagate a smoother demand stream by requesting a sequence of deliveries that is evenly spaced over time. Such smoothing might reduce cycle stock at the stage in question, reduce safety-stock requirements of upstream stages, and offer other benefits such as even production loading. This section extends the GSM framework with a simple version of smoothing.

Although one might imagine myriad variations of smoothing, our simple approach requires just one additional input parameter that we refer to as η, the replenishments-per-review period. We assume that at each review period, a stage calculates a total order quantity as in Equation (4), but divides this quantity into η equally sized fractional orders due at equally spaced intervals, with the first due a service time from the current review period. Consequently, we require that η be a factor of R for each stage; we can relax this factor assumption somewhat, at the expense of additional algebra. We assume that the upstream stage is largely unaware of the smoothing. More specifically, although the supplier knows the timing of orders, it does not or cannot exploit the equality in size of η consecutive orders. It essentially acts as if its customer has a review period of R/η.

The inventory balance equation is now

$$\begin{aligned} I(t) = {} & B(t - SI - T - x'(t) - (R-1)) \\ & - d(t - SI - T - (R-1) - x'(t), t - S) \\ & + \frac{\eta(t)}{\eta} \cdot d(t - SI - T - (R-1) - x'(t), \\ & \qquad t - SI - T - x'(t) + 1) \\ & + \frac{\eta(t)}{\eta} \cdot [B(t - SI - T - x'(t) + 1) \\ & \qquad - B(t - SI - T - (R-1) - x'(t))], \end{aligned} \tag{5}$$

where $\eta(t) = \lceil x'(t)/(R/\eta) \rceil$ and $x'(t) = (t - SI - T - \omega - (R-1)) \bmod R$. The first term is the base-stock target of the most recent review period whose sequence of fractional orders has been completely supplied. The second term functions similarly to the inventory-exposure term of Equation (1); it includes demands that are due but are not yet fully replenished. Its first argument, $t - SI - T - (R-1) - x'(t)$, corresponds to the former review period, and the latter, $t - S$, to the most recent demand that is due. The third and fourth terms of Equation (5) correspond to fractional orders already received from the current incoming sequence. The numerator, $\eta(t)$, is the number of such fractions already received, and the ratio $\eta(t)/\eta = 1$ for the first time in an inventory cycle when $x'(t) = ((\eta - 1)/\eta) \cdot R + 1$, on arrival of the ηth and final fraction. The second factors of the third and fourth terms, those multiplied by $\eta(t)/\eta$, define the bulk order calculated per Equation (4) for the current incoming sequence, and so consist of demand incurred over the corresponding review cycle and a base-stock adjustment.

Application at Celanese

Celanese is the world's largest producer of acetyl products, including acetic acid monomer, vinyl acetate monomer, and polyacetal products. It is also a leading global producer of high-performance, engineered polymers for consumer and industrial applications. A key factor in Celanese's success has been its hybrid business model of producing chemical building blocks, derivative products, and engineered finished products, instead of focusing on one market. The company, which had net sales of \$6 billion in 2005, is organized into four business segments.

The end-to-end supply chain at Celanese is complex and vast. It stretches from natural gas raw materials to basic chemicals to advanced polymers. If the supply chain is divided on the basis of the dimension of business control, it resembles three production–distribution networks, each serving global demand for its own products as well as internal processes that utilize the product as the main material source for a totally different downstream product. For instance, acetic acid is used to produce vinyl acetate; in turn, vinyl acetate is the main ingredient of polyvinyl alcohol (PVOH). This end-to-end supply chain has over 3,000 stages if we consider all the late-stage differentiation for PVOH. Within this supply chain, material typically moves in large quantities because of economies of scale and transportation schedules. Although the actual production times are smaller than transportation times in the supply chain, distribution of goods in a global environment exposes the reality that transportation modes have defined schedules of operation. This intermittent behavior is also apparent in customer orders: they are typically batched and timed with weekly, biweekly, or monthly frequencies. Looking at the processing times independent of the ordering frequency would not accurately represent the inventory dynamics of the supply chain.

In this section, we illustrate the significance of our review-periods functionality to appropriately modeling the dynamics of Celanese's supply chain. We begin by describing the acetic acid supply chain; acetic acid is the primary product offering of the Chemical Products business segment, which is responsible for 71 percent of Celanese's 2005 net sales. We describe the chain as a whole, as modeled in *PowerChain Inventory*, and then we analyze a stylized reduction of the chain in more depth. This section includes two figures from the *PowerChain Inventory* user interface. A square box or icon corresponds to a stage (recall that an arc defines only a precedence relationship between stages). In addition, a triangle next to a stage indicates the presence of safety stock. By default, the review period at each stage is 1, a single base time unit. A user may "enable" review-periods functionality on a stage-by-stage basis, and for each stage so activated, enter a review period (R), offset (ω), and replenishments per review-period parameter (η).

Acetic Acid Supply Chain

Acetic acid is a building block for a variety of industrial products, including colorants, paints, adhesives, coatings, plastics, medicines, cosmetics, detergents, textiles, and fragrances. Figure 3 models the entire end-to-end supply chain.

Stages in the supply chain correspond to vendors supplying a liner used to transport the acetic acid, manufacturing sites producing acetic acid, transportation modes moving product, warehouses storing acid, and geographically specific demand locations. There are four manufacturing sites. U.S. locations are in Pampa and Clear Lake in Texas and Calvert City in Kentucky; there is also one plant in Singapore. These plants supply several echelons of storage locations in Asia, Europe, and the Americas. Transportation stages correspond to rail, barges, trucks, and ocean vessels. A downstream storage facility is commonly served by multiple upstream facilities, as well as directly from a manufacturing site.

There are 90 stages and 94 links in the supply chain. The base time unit is one day. The longest stage

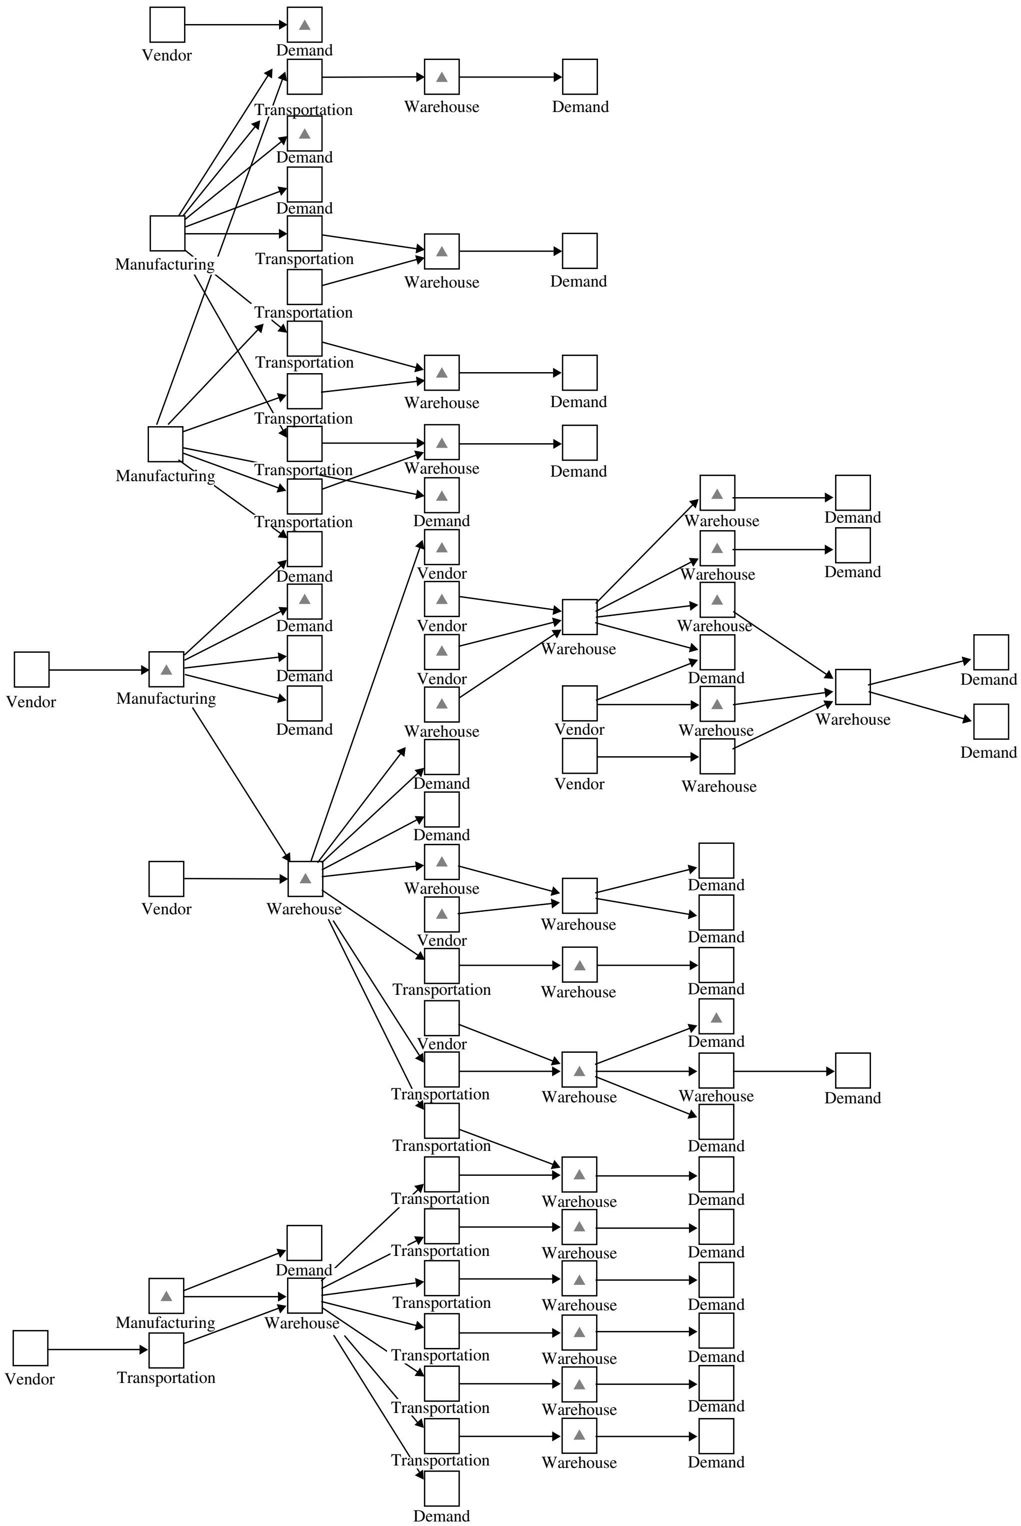

Figure 3: The graphic shows the end-to-end Celanese acetic acid supply chain as modeled in *PowerChain Inventory*.

time path through the network is 50 days. Of the 90 stages, 34 enable review periods. Enabled review periods range from three to 30 days; review-period offsets range from two to 26 days. Review periods are associated with all of the storage locations and some of the demand locations. A review period at a storage location reflects the fixed schedule of the supplying transportation mode.

Analysis of Stylized Chain

We have extracted a subset of the stages in the supply chain to demonstrate the significance of review periods more clearly (Figure 4).

To protect company confidentiality, we have disguised the data in the supply chain; in addition, reducing the supply chain from 90 to 14 stages eliminates some detail. In particular, we aggregated demand at each location into one demand class; in reality, there can be multiple demand classes segmented by customer and by use. Each manufacturing and storage location is now sole sourced. However, the resulting chain still reflects Celanese's business.

The supplier stage provides the transportation liner to the manufacturing site that produces acetic acid. The manufacturing site serves some demand directly and also supplies the central warehouse, which ships to two customer regions by boat. Region 3 supplies an additional storage location that serves the region's demand. The port warehouse in region 1 supplies region 2's warehouse by sea plus its own satellite warehouses by rail and truck. There are five demand stages—one for each of the regional warehouses and one for demands that are placed directly on the manufacturing site. These five stages characterize the demand-generation process rather than a physical operation, so their stage costs are zero, stage times are zero, and they quote service times of zero.

Daily demand at each demand stage has mean and standard deviation of 500,000 metric tons. The target service level at each demand stage is 95 percent, and the annual holding-cost rate at each stage is 25 percent. Each stage has only one replenishment per review period and employs the constant safety-stock model, which mimics Celanese's actual business practice.

The review periods do not nest or follow any other apparent structure. The warehouses supplied by region 1's port warehouse, review more frequently than its port warehouse, whereas the regional warehouse in region 3 reviews less frequently than its port warehouse. The supplier to manufacturing reviews every three days, and manufacturing itself reviews every day.

To assess the importance of correctly modeling review periods, we present numerical results for two models. The review-periods model (RPM) employs the data in Table 2 and the modeling approach we outlined in the *Extension for General Review Periods* section to determine the inventory levels and locations across the supply chain. The consolidated time model (CTM) ignores the distinct effects of review periods. For each stage, the CTM simply adds the review period minus one to the stage time in Table 2 to produce a new stage time. The base time unit is assumed to be the common review period of all stages, and the algorithm from GW is run.

Table 3 summarizes the inventory metrics for the RPM, and the triangles of Figure 4 indicate the corresponding inventory locations. Table 4 summarizes the inventory metrics for the CTM; the CTM's calculated stage time is also included as an extra column in Table 4.

Inventory levels from the CTM differ significantly from those of the RPM. The RPM has 51 percent less safety stock than the CTM, but its on-hand inventory of cycle and safety stock is 32 percent greater than that

Stage name	Stage cost ($)	Stage time	Review period	Review-periods offset
Central warehouse	0.05	1	1	0
Demand—manufacturing	0.01	0	1	0
Demand Region 1A	0.01	0	1	0
Demand Region 1B	0.01	0	1	0
Demand Region 2	0.01	0	1	0
Demand Region 3	0.01	0	1	0
Manufacturing	0.25	5	1	0
Region 3 regional warehouse	0.05	2	14	3
Ship Region 1 port warehouse	0.25	19	14	10
Ship Region 2 warehouse	0.05	3	3	0
Ship Region 3 port warehouse	0.20	7	10	7
Supplier	0.50	5	3	0
Warehouse Region 1A	0.05	3	7	4
Warehouse Region 1B	0.05	4	7	3

Table 2: The data in the table show the disguised cost and lead-time data of the stylized chain.

Stage name	Safety-stock inventory	Cycle-stock inventory	Pipeline-stock inventory	Service time	Net-replenishment lead time
Central warehouse	—	—	2,000,000	1	0
Demand—manufacturing	—	—	—	0	0
Demand Region 1A	—	—	—	0	0
Demand Region 1B	—	—	—	0	0
Demand Region 2	—	—	—	0	0
Demand Region 3	—	—	—	0	0
Manufacturing	6,632,384	—	12,500,000	0	12
Region 3 regional warehouse	2,254,464	3,250,000	1,000,000	0	32
Ship Region 1 port warehouse	3,488,128	8,250,000	28,500,000	0	33
Ship Region 2 warehouse	1,346,528	500,000	1,500,000	0	5
Ship Region 3 port warehouse	—	2,500,000	3,500,000	17	0
Supplier	—	2,500,000	12,500,000	7	0
Warehouse Region 1A	1,296,128	1,500,000	1,500,000	0	9
Warehouse Region 1B	1,406,720	1,500,000	2,000,000	0	10

Table 3: We show inventory levels by stage with review periods enabled.

of the CTM, which has no cycle stock. Finally, the sum of on-hand inventory and pipeline stock is 30 percent higher in the CTM than the RPM. We expect these qualitative results to hold in general. Because review periods increase upstream demand variability and introduce cycle stock, which contributes less directly to service than safety stock, a model that accurately represents their dynamics should have greater on-hand inventory than a model that simply consolidates time. On the other hand, the latter approach will quote excessive pipeline stock, which is proportional to stage times and mean demands.

Although the above inventory-level discrepancies stem primarily from the different single-stage inventory dynamics of the models, the discrepancies can lead to different supply chain stocking strategies. For all stages including and downstream of the central warehouse stage (CW), both models report the same optimal locations, albeit with different inventory values. The port warehouse in region 1 holds a

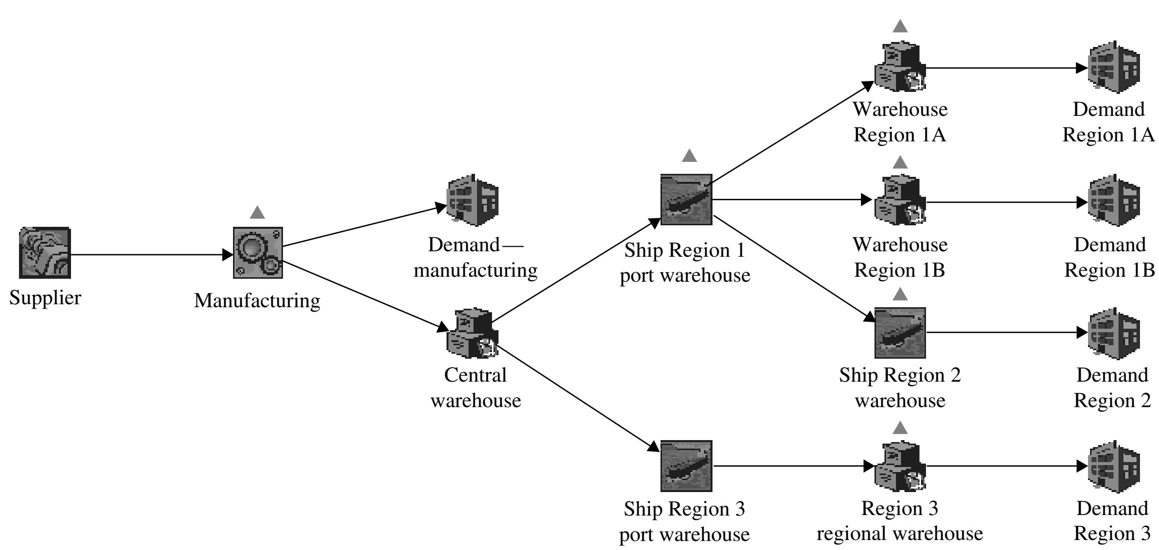

Figure 4: We show a 14-stage version of the acetic acid supply chain.

Stage name	Safety-stock inventory	Cycle-stock inventory	Pipeline-stock inventory	Service time	Net-replenishment lead time	Stage time
Central warehouse	—	—	2,000,000	13	0	1
Demand—manufacturing	2,848,970	—	—	0	12	0
Demand Region 1A	—	—	—	0	0	0
Demand Region 1B	—	—	—	0	0	0
Demand Region 2	—	—	—	0	0	0
Demand Region 3	—	—	—	0	0	0
Manufacturing	—	—	12,500,000	12	0	5
Region 3 regional warehouse	5,455,362	—	7,500,000	0	44	15
Ship Region 1 port warehouse	9,555,736	—	48,000,000	0	45	32
Ship Region 2 warehouse	1,839,002	—	2,500,000	0	5	5
Ship Region 3 port warehouse	—	—	8,000,000	29	0	16
Supplier	—	—	17,500,000	7	0	7
Warehouse Region 1A	2,467,280	—	4,500,000	0	9	9
Warehouse Region 1B	2,600,742	—	5,000,000	0	10	10

Table 4: We show inventory levels by stage when review periods are simply added to stage times.

decoupling safety stock that pools the demand variability from regions 1 and 2. The port warehouse in region 3 does not hold safety stock. Instead, the regional warehouse pools the demand variability over the time that elapses between the CW and region 3 regional warehouse.

The policies upstream of the CW differ. For the RPM, manufacturing holds a decoupling safety stock sufficient to buffer demand from all regions plus its direct demand, over the review period and stage times that are associated with both the supplier and manufacturing stages. In contrast, for the CTM, the safety stock at manufacturing addresses only the variability of its direct end-item demand. In turn, the decoupling stocks at the region 1 port warehouse and the region 3 regional warehouse cover stage times from the supplier and manufacturing stages, as well as the CW. For the CTM, these two warehouse stages would have significant net-replenishment lead times even if manufacturing quotes a service time of zero. Because the majority of cost accrues before manufacturing, increasing safety stock at the warehouse stages is more economical than forcing manufacturing to cover all the demand variability in the supply chain, even over its relatively short net-replenishment lead time.

In general, we have seen that chains that have review periods enabled tend to hold more inventory upstream in the supply chain, as in the stylized Celanese chain. If the changes in both demand variability and timing are not rigorously justified, it is often beneficial to hold no upstream inventory and effectively add the lead time from the upstream stages to downstream stages that are already pooling across time and demand locations.

An alert modeler might readily address shortcomings of the CTM, but basic observations would likely go only so far. Because pipeline stock is directly proportional to stage time, one could correct the CTM excess through the ratio of the true stage time to the CTM-inflated stage time. On the other hand, interactions among safety stock, cycle stock, and input model parameters are isolated less easily. Even if one could infer a reasonable cycle-stock value, the appropriate reduction of safety stock would remain unclear because the relationship between cycle stock and service is unclear. Indeed, complications of this nature led to our use of search algorithms for base-stock targets. Furthermore, average cycle stock is not, in general, half the demand over a review cycle. Rather, it depends on the timing of outgoing demands and incoming replenishment. For example, a stage receiving intermittent demand can avoid cycle stock entirely if its replenishments coincide with demand due dates.

To this point, we note that the review-periods model adds a new dimension to the optimization—that of coordinating service times to reduce cycle stock; however, the CTM is blind to cycle stock. For the 14-stage chain we considered here, the CTM outputs a very reasonable policy; plugging its service times into the review-periods model increases on-hand cost by only 3.87 percent. However, an example of just two stages can illustrate that such policy quality does not necessarily hold.

Conclusion

We incorporate stage-dependent review periods in the guaranteed service, supply chain modeling framework in a manner that addresses numerous real-world complications. We do not constrain the review periods across a supply chain to a well-behaved structure or pattern. In addition, our framework accommodates a variety of operating policies, including constant base-stock targets, constant safety-stock targets, and adaptive base-stock targets, as well as a variation of demand smoothing. As a form of validation, we analyze the Celanese acetic acid supply chain, which entails some of these complications. The review-periods model yields inventory metrics that differ by more than 30 percent relative to the simpler approach of adding the review period to the stage time.

In practice, we have extended the presentation here to several other real-world complications, including general acyclic networks, stochastic lead times, and time-phased demand. We are currently developing extensions to alternative service metrics and batch ordering. Although these capabilities enhance applicability of the modeling framework, we excluded them from this paper to focus on the problems that review periods pose.

We close by noting several remaining questions. First, the smoothing extension of the *Demand Smoothing* section is just one of many possible models of this potentially complicated behavior. This capability remains new, and we await feedback on its general applicability. In addition, we suspect that more refined approximation schemes could further improve the accuracy of optimization under review periods. In particular, chains with adaptive policies and unstructured review periods remain sensitive to the circularity issue noted in the *Optimization* section. Finally, the computational expense of evaluating a chain grows with the cycle lengths λ^{in} and λ^{out}, and, in turn, large chains or chains without nested review periods can become computationally burdensome. Evaluation of a candidate base-stock target currently entails calculating the service level at each period of the cycle. We have begun researching several approximation schemes that show initial promise for streamlining the calculations.

Appendix 1

Reviewing the GSM Framework

This section presents an overview of the modeling assumptions, inventory dynamics, and mathematical programming formulation of the GSM framework. The formulation in this section follows the notation in Graves and Willems (2000).

GW assume that the supply chain can be represented as a network with node set N and arc set A. Each stage in N corresponds to a process at the end of which we might hold inventory, and the arc set A defines the precedence relationships among stages. In particular, a transportation process is modeled as a stage, not an arc. A stage does not necessarily correspond to a specific facility within the physical supply chain. Depending on the level of granularity required, individual stages might represent intermediate inventories within a single plant, or a single stage might justifiably aggregate several physical locations.

Each stage $j \in N$ has a deterministic processing time T_j and follows a periodic-review, constant base-stock policy. All stages have the same review period that equals the base time unit of the model. External demand originates only at stages with no outgoing arcs; external demand arrives each base time unit, and the demands over time are stationary and independent. Under a base-stock policy, at each review period a stage orders to reset its inventory position, which includes both on-hand and on-order inventory, to a precalculated base-stock target. When the targets are constant, each stage simply orders the demand placed on it since it last reviewed, and external demand thereby immediately propagates across the chain, although it might be inflated through bill-of-materials (BOM) multipliers. Specifically, $d_j(t)$, the demand received by stage j over time t, is calculated as $d_j(t) = \sum_{k \mid (j,k) \in A} \phi_{jk} d_k(t)$, where ϕ_{jk} is the number of units from stage j required to produce one unit of output at stage k.

Perhaps most critically, GW also assume guaranteed service and bounded demand. Demand is bounded at each stage j in the sense that for any relevant length of time τ (and arbitrary time t), there is a known bound $D_j(\tau)$ such that $D_j(\tau) \geq \sum_{s=1}^{\tau} d_j(t+s)$. Next, each stage j quotes an outbound service time S_j to its immediate customers such that it will deliver exactly $d_j(t)$ at time $t+S_j$. GW and Graves and Willems (2003) discuss these key assumptions in much greater detail.

GW balance the stage j net inventory level as $I_j(t) = B_j - d_j(t - SI_j - T_j, t - S_j)$, where B_j is the stage j base-stock target, SI_i is its inbound service time, and $d_j(t_1, t_2) = \sum_{s=t_1+1}^{t_2} d_j(s)$. The order stage j places at time t arrives in inventory at $t + SI_j + T_j$. Equivalently, the last replenishment received by time t was placed at $t - SI_j - T_j$. Similarly, the last demand served by stage j by time t was placed at $t - S_j$. Therefore, at each time, the inventory is exposed to $SI_j + T_j - S_j$ periods of demand. This constant exposure length is referred to as the *net-replenishment lead time*. To achieve guaranteed service with minimum inventory, GW set the base stock to the demand bound corresponding to the net-replenishment lead time, that is, $B_j = D_j(SI_j + T_j - S_j)$. To model a supply chain given guaranteed service times S_j, they link stage-inventory balances by propagating demand upstream through the arc multipliers ϕ_{jk}, determining a demand-bounding function $D_j(\tau)$ for each stage j, and calculating an inbound service time SI_j to each stage j as $SI_j = \max_{i\,|\,(i,j)\in A}\{S_i\}$.

GW propose a deterministic, dynamic program to minimize the holding cost incurred across a supply chain that is modeled within the GSM framework as a spanning tree. Specifically, they solve problem (P):

$$\text{(P)} \quad \min \sum_{j=1}^{|N|} h_j \cdot [D_j(SI_j + T_j - S_j) - (SI_j + T_j - S_j) \cdot \mu_j]$$

$$\begin{aligned} \text{s.t.}\quad & S_j - SI_j \leq T_j \quad \forall j \in N, \\ & SI_j - S_i \geq 0 \quad \forall (i,j) \in A, \\ & S_j \leq s_j \quad \forall j\colon \exists k \in N \mid (j,k) \in A, \\ & S_j, SI_j \geq 0, \text{ integral} \quad \forall j \in N, \end{aligned}$$

where h_j is the stage j holding cost, s_j is a bound on the outbound service time that a demand stage j can quote, and μ_j is the mean demand received per period by stage j. The decision variables are the service times S_j and SI_j, and they affect inventory levels and locations. The first set of constraints ensures nonnegative net-replenishment lead times, and the second set enforces the definition of inbound service time. GW describe a functional equation for the cost over a subtree in terms of the optimal costs for the subtrees connected to a single stage. Each subproblem of their dynamic program optimizes over a single service time by iteratively evaluating this equation, and the dynamic program entails $|N|$ subproblems.

Appendix 2

Computational Study of Five-Stage Serial Supply Chains

This appendix summarizes a computational study on the accuracy of the optimization approximations described in the *Extension for General Review Periods* section. We consider 189 permutations of the five-stage serial-line supply chain that Figure A.1 depicts.

End-item demand originates at stage 5. Daily demand at stage 5 is characterized as a normally distributed random variable with average and standard deviation equal to 100. Across all permutations, the service level at stage 5 is 95 percent, the holding-cost rate at all stages is 35 percent, and each stage has one replenishment per period.

The 189 chains represent the permutations of three cost-accrual profiles, three lead-time profiles, and 21 review-period profiles. The direct costs added per unit at stages 1 through 5 under the three cost profiles are (\$10, \$4, \$3, \$2, \$1), (\$4, \$4, \$4, \$4, \$4), and (\$1, \$2, \$3, \$4, \$10). Similarly, the three time profiles, in days, are (10, 4, 3, 2, 1), (4, 4, 4, 4, 4), and (1, 2, 3, 4, 10). The 21 review-period profiles, also in days, comprise seven sets of three profiles: starting, middle, ending, uniform, increasing, decreasing, and random. The starting profiles [(2, 1, 1, 1, 1), (5, 1, 1, 1, 1), (10, 1, 1, 1, 1)], middle profiles

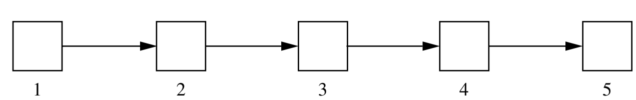

Figure A.1: Five-stage serial-line supply chain used for numerical analysis.

[(1, 1, 2, 1, 1), (1, 1, 5, 1, 1), (1, 1, 10, 1, 1)], and ending profiles [(1, 1, 1, 1, 2), (1, 1, 1, 1, 5), (1, 1, 1, 1, 10)] have a single enabled review period. The uniform profiles [(2, 2, 2, 2, 2), (5, 5, 5, 5, 5), (10, 10, 10, 10, 10)] maintain the same review period across the supply chain. The review periods of increasing profiles increase upstream: [(16, 4, 4, 4, 2), (10, 8, 6, 4, 2), (14, 7, 6, 3, 2)]. Similarly, the decreasing profiles are [(2, 4, 4, 4, 16), (2, 4, 6, 8, 10), (2, 3, 6, 7, 14)]. The random profiles [(2, 4, 16, 4, 4), (10, 6, 4, 8, 2), (7, 14, 3, 2, 6)] follow no predictable pattern.

In the presence of review periods, the *Extension for General Review Periods* section defines the operating characteristics and optimization problem for supply chains operating a constant base-stock policy, an adaptive base-stock policy, and a constant safety-stock policy. Table A.1 compares the performance of the optimization approach we described in the *Extension for General Review Periods* section against exhaustive enumeration for each of the three policies.

We presented the first column of Table A.1 in the *Optimization* section. We express the errors in terms of the expected on-hand inventory cost, the sum of expected safety stock, and expected cycle-stock costs. For a single permutation, we calculate the approximation error as the absolute difference between the exact solution and the approximate solution divided by the exact solution. The conditional average error reflects only those permutations with positive errors.

Recall that a stage under a constant base-stock policy orders at each review period demand observed since it last reviewed. Consequently, such chains are not subject to the circularity between service times

	Constant base-stock policy	Adaptive base-stock policy	Constant safety-stock policy
Average approximation error (%)	1.27	3.01	6.87
Maximum approximation error (%)	16.82	76.49	55.87
Number of permutations where error is zero	131	92	92
Number of permutations with nonzero error	58	97	97
Conditional average approximation error (%)	4.14	5.86	13.39

Table A.1: We show summary performance metrics for each of the three inventory policies we presented in the *Extension for General Review Periods* section.

Review-periods profile	Constant base-stock policy (%)	Adaptive base-stock policy (%)	Constant safety-stock policy (%)
Start	0.42	0.29	0.29
Middle	0.28	0.60	0.30
End	0.19	1.02	1.66
Uniform	0.11	0.23	0.23
Increasing	2.40	3.47	4.88
Decreasing	3.68	4.47	21.92
Random	1.81	10.99	18.82

Table A.2: We show the average approximation error across all permutations that correspond to the associated review-periods profile.

and propagated demand that we noted in the *Extension for General Review Periods* section, and their relatively superior performance is not surprising. We do not yet understand the markedly better accuracy under fully adaptive policies versus constant safety-stock policies, although we can attribute the performance gap primarily to the decreasing and random review-periods profiles. Table A.2 disaggregates the average error results by review-period profile; similar disaggregations over cost and time profiles reveal no obvious trends.

Acknowledgments

We thank Stephen C. Graves (Massachusetts Institute of Technology) and John Neale (Optiant, Inc.) for their helpful and detailed comments on an earlier draft of this paper. We gratefully acknowledge the computational assistance of Salal Humair (Optiant, Inc.) with the numerical study we report in Appendix 2.

References

Billington, C., G. Callioni, B. Crane, J. D. Ruark, J. U. Rapp, T. White, S. P. Willems. 2004. Accelerating the profitability of Hewlett-Packard's supply chains. *Interfaces* **34**(1) 59–72.

Federgruen, A. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. S. C. Graves, A. H. Rinnooy Kan, P. H. Zipkin, eds. *Handbooks in Operations Research and Management Science*, Vol. 4. *Logistics of Production and Inventory*. North-Holland Publishing Company, Amsterdam, The Netherlands, 133–173.

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

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

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 Operations Research and Management Science*, Vol. 11. *Supply Chain Management: Design, Coordination and Operation*. North-Holland Publishing Company, Amsterdam, The Netherlands, 95–132.

Hadley, G., T. M. Whitin. 1963. *Analysis of Inventory Systems*. Prentice Hall, Englewood Cliffs, NJ.

Humair, S., S. P. Willems. 2006. Optimal inventory placement in networks with clusters of commonality. *Oper. Res.* **54**(4) 725–742.

Lee, H. L., C. Billington. 1993. Material management in decentralized supply chains. *Oper. Res.* **41**(5) 835–847.

Simpson, K. F. 1958. In-process inventories. *Oper. Res.* **6**(6) 863–873.

van Houtum, G. J., A. Scheller-Wolf, J. Yi. 2003. Optimal control of serial, multi-echelon inventory/production systems with periodic batching. Working paper, Technische Universiteit Eindhoven, Eindhoven, The Netherlands.

Todd Carter, director of global supply chain at Celanese, writes: “Celanese has been a Six Sigma company since 1997 and supply chain has been Six Sigma since 2002. We are a data driven company and we needed a tool that understood that. Our old inventory optimization planning consisted of linked Excel spreadsheets and the Solver tool. The results were not always accurate and it was very hard to update actual targets. Furthermore, our old tool used monthly data instead of daily data.

“We knew already through our Six Sigma work that using daily numbers was more important. We needed to set our inventory targets more frequently; therefore, we needed a tool to reduce the cycle time of setting inventory targets. Finally, we had never proved to the business that our targets were optimal. Optiant PowerChain allows us to do this.

“We have built an Optiant model for every molecule that we sell. It is a global model, understanding lead times, safety stock, points of manufacture, points of consumption, and planned service levels. The tool provides us with a fully costed supply chain. We are able to see where inventory and its costs accumulate. We also understood our lead time and lead time variabilities across our manufacturing and transportation lanes. We are also able to model our demand variability by region. This would not have been possible if we could not exactly model the operational limits of our supply chain. In particular, different locations have different capacity to hold inventory and these different locations have different restocking frequencies.

“We have had a significant reduction in inventory using Optiant. Prior to using Optiant, we were already on an inventory decline as a function of sales. Optiant accelerated that decline. We also and very importantly were able to simultaneously improve fill rates. That is, we were able to enter planned service levels in Optiant and we then met those fill rates when we implemented the targets.

“Because Optiant’s tool matches the realities of the chemicals business, it has become an accepted way to make supply chain decisions in Celanese.”