Coordination of Installation Base-Stock Policies in Supply Chains with Compound Poisson Demand

Yao Zhao

Department of Management Science and Information Systems

Rutgers University, Newark, NJ

1st Version, June 2005; Revised, December 2005

Abstract

We consider a supply chain with tree network structure where external demands follow independent compound Poisson processes and each stage controls the inventory of one product by an installation base-stock policy. The inventories in the supply chain are either reviewed continuously or periodically in time. The lead-times are stochastic and sequential. Unsatisfied demands at each stage are fully backordered. We characterize the backorder (or stock-out) delay for each unit of a demand at each stage of the supply chain, and present an exact and systematic approach to analyze various material flow topologies in tree networks. For supply chains under continuous-review base-stock policies, we demonstrate the similarities and structural differences between compound Poisson demand and Poisson demand. We also compare and contrast the supply chains under continuous-review base-stock policies with those under periodic-review base-stock policies. Based on the analyzes, we present simple and tractable approximations which facilitate efficient coordination of the installation inventory policies at all stages with the objective of minimizing the system-wide inventory cost subject to certain service requirements of the external customers. We demonstrate the effectiveness of the coordination by numerical studies.

Key words: Evaluation and coordination, installation inventory policies, compound Poisson demand, continuous-review, periodic-review, stochastic sequential lead-time, tree network

1 Introduction

The control of inventories in complex network of production and distribution facilities poses a substantial challenge to many manufacturing and logistics firms. Because of the inherent uncertainties in demand, manufacturing and transportation processes, companies often keep millions of dollars of inventory in their supply networks in order to provide satisfactory service to external customers (Feigin 1999).

Due to the complexities of many real world supply chains and the internal organizational barriers, simple and easily implementable installation inventory policies are often employed in practices. Installation inventory policies require minimum information transferring among different facilities in a firm, and more importantly, they allow each facility to manage its inventory autonomously. To execute an installation policy, a facility only needs to know its local demand processes (from immediate downstreams) and order status (from immediate upstreams). Examples can be found in many industries, such as computer (Lin, et al. 2000), electronics (Lee and Billington 1993) and consumer packaged products (Graves and Willems 2000). Studies of these examples demonstrate that proper coordination of installation policies at different facilities could result in significant inventory cost reduction and customer service improvement.

In this paper, we consider a class of supply chains with tree structure where external demands follow independent compound Poisson processes. The processing cycle times and transportation lead-times are stochastic, sequential and exogenous. Each stage controls its inventory by either a continuous-review or a periodic-review installation base-stock policy. Our objective is to characterize the performance metrics of these supply chains, and to develop simple and tractable approximations which allow for efficient and effective coordination of the installation policies at all stages of the network.

For continuous-review supply chains with compound Poisson demand, we characterize the stockout delay for each unit in a demand at each stage of the supply chain, and provide a unified and exact approach to analyze various network topologies, such as serial, pure assembly, pure distribution and assembly-distribution systems (see Sections 3.3-3.5 for definitions). We then demonstrate the similarities and structural differences between supply chains facing compound Poisson demand and those facing Poisson demand. Based on the exact analysis, we present simple approximations which serve as a basis for system optimization subject to type 2 fill rate constraints within certain committed service times. The similarities and differences between continuous-review supply chains and periodic-review supply chains are characterized, and the analysis and approximations are extended to the supply chains where each stage employs a periodic-review base-stock policy. Numerical studies are conducted to identify the conditions under which the approximations are reasonably accurate, and to demonstrate the effectiveness of the solutions.

The continuous-review model can be applied to low volume but high risk items such as those carried by service part supply chains, where batch demand arrivals are common and the inventories are managed by continuous time base-stock policies. We refer the reader to Sherbrooke (1968), Muckstadt (1973) and Graves (1985) for real world examples. Muckstadt (2004) provides an excellent overview. The periodic-review model can be applied to high volume items carried by supply chains where inventories are reviewed periodically and are managed by base-stock policies. Real world examples can be found in Lee and Billington (1993) and Graves and Willems (2000).

This paper is organized as follows: we review the related literature in Section 2. In Section 3, we characterizes the performance metrics of various network topologies for the continuous-review supply chains; while in Section 4, we present the analysis for the periodic-review supply chains. Approximations are presented and optimization problems are formulated in Section 5. Numerical studies are conducted in Section 6. Finally, we conclude the paper in Section 7.

2 Literature Review

Inventory control in multi-stage supply chains has been studied extensively over the past forty years. The related literature to this paper is voluminous. We do not attempt to provide an exhaustive review of all the related work, rather, we focus on most related papers and refer the reader to Federgruen (1993), Porteus (2002) and Zipkin (2000) for excellent reviews.

For serial systems with stochastic demand, the optimal policies are characterized by Clark and Scarf (1960), Muharremoglu and Tsitsiklis (2002) and Parker and Kapuscinski (2004) in various settings. For assembly systems with stochastic demand but constant lead-times, we refer to Schmidt and Nahmias (1985) for insights on the structure of the optimal policy, to Rosling (1989) for the characterization of optimal policy in multi-stage assembly system, and to Chen (2000) for the analysis of batch ordering assembly systems.

For more general supply chains, such as a tree structured network with random demand and stochastic lead-times, the optimal inventory policy remains unknown. An alternative approach is policy evaluation and optimization, that is, given a certain inventory policy for each stage, evaluate and/or optimize the performance of the entire system. There is a rich literature in this area. For serial and distribution systems with installation inventory policies, we refer the reader to the following work: Sherbrooke (1968, 1986), Muckstadt (1973), Deuermeyer and Schwartz (1981), Graves (1985, 1996), Zipkin (1991), Lee and Moinzadeh (1987a, b), Sovoronos and Zipkin (1988, 1991), Glasserman and Tayur (1995, 1996), Axsater (1993a, 2000). For serial and distribution systems with echelon inventory policies, we refer the reader to Chen and Zheng (1994), Chen and Zheng (1997) and Gallego and Zipkin (1999). Axsater (2002) provides an excellent survey. For assembly systems with non-zero lead-times, we refer the reader to Song (1998, 2002), Song and Yao (2000), Lu and Song (2004) for Poisson demand and i.i.d. lead-times, Zhao and Simchi-Levi (2005) for Poisson demand and stochastic sequential lead-times, and finally, Song (2000) for models with constant lead-times and compound Poisson demand. Song and Zipkin (2002) provides an excellent literature review.

A few recent work developed unified approaches to analyze supply chains with more general network structures which include serial, assembly and distribution systems, e.g., Lee and Billington (1993), Graves and Willems (2000), Ettl, et al. (2000) and Simchi-Levi and Zhao (2005). These work are most related to this paper. Detailed reviews of these papers are provided by Graves and Willems (2002) and Simchi-Levi and Zhao (2005). Here, we focus on the differences and similarities between these work and this paper.

One key element of these work is to characterize or determine the delays due to stockout at each stage of the supply chain. Lee and Billington (1993) considers periodic review supply chains and derive simple approximations for the delays due to stockout at each stage of the network. The focus of Lee and Billington (1993) is on performance evaluation. In contrast, this paper provides exact analysis and focuses on system optimization. Graves and Willems (2000) considers similar supply chains with tree structure under the guaranteed service assumption, that is, if a demand cannot be satisfied within the committed lead-time at a stage due to stockout, the stage has resource other than the on-hand inventory to satisfy the demand. Therefore, the service is 100% guaranteed at all internal stages and the lead-times are constants. In this paper, we assume that unsatisfied demands are backordered until the on-hand inventory becomes available (full backorders).

Assuming full backorders, Ettl, et al. (2000) develops a model to evaluate and optimize supply chains with compound Poisson demand and continuous-review base-stock policies. While Ettl, et al. (2000) models the lead-times by i.i.d random variables, we model the lead-times by stochastic sequential random variables. The model of stochastic sequential lead-times, i.e., the *transit time*, is formally proposed by Sovoronos and Zipkin (1991). In this model, the processing cycle times and transportation lead-times are mutually independent and are also independent of the system state. Furthermore, consecutive orders cannot cross over, and hence Palm's theorem (Palm 1938) cannot be applied. Sovoronos and Zipkin (1991) points out that this lead-time model may be more realistic than the i.i.d. lead-time model in some real world applications, see Zipkin (2000) for more discussions. Simchi-Levi and Zhao (2005) considers tree structure supply chains with stochastic sequential lead-times where each stage controls its inventory by a continuous-review base-stock policy. For point demand processes, the authors derived sample path based recursive equation for the backorder delay at each stage of the supply chain; while for independent Poisson demands, they characterized the performance of various network topologies.

As Simchi-Levi and Zhao (2005) point out, compound Poisson demand significantly complicates the probabilistic analysis of tree structure supply chains, partly because different units in the same demand face statistically different backorder delays at each stage, partly because the supply chains facing compound Poisson demands have different dynamics than those facing Poisson demands; see also Zipkin (1991, page 405). Exact analysis and/or approximations are developed for various serial and distribution systems facing compound Poisson demand, see Graves (1985), Zipkin (1991), Axsater (2000) and reference therein. So far, exact analysis of general assembly-like systems with stochastic sequential lead-times and compound Poisson demand is not available. We also lack of a systematic and exact approach that can handle all material flow topologies in the more general tree networks. Furthermore, approximations and efficient algorithms need to be developed to compute the optimal or near optimal stock levels at all stages so as to minimize the system-wide inventory cost subject to certain service level requirements of the external customers. We refer the reader to Axsater (2002) and Section 5 for the importance of approximation techniques in optimizing "larger" size problems.

This paper takes one step in filling these gaps. Its main contribution is providing an exact and systematic approach to analyze tree structure supply chains facing compound Poisson demands, where the lead-times are stochastic and sequential, and each stage controls its inventory by either a continuous-review bases-stock policy or a periodic review base-stock policy. The approach serves as a basis for comparing and contrasting *analytically* various supply chains with either compound Poisson demand or Poisson demand, under either continuous-review or periodic-review. Finally, the exact analysis allows us to develop and test tractable approximations, which lead to efficient coordination of relatively large systems (see Section 6).

3 The Continuous-Review Supply Chains

In this section, we consider the tree structure supply chains where each stage manages one product and controls its inventory by a continuous-time base-stock policy with a non-negative base-stock level. Each stage in the supply chain consists of a processing facility and a storage facility. It could be a store, a distribution center or a manufacturing plant. The processing cycle time at each stage (i.e., the time between item inception to finishing production), the transportation lead-time between every two stages and the lead-times from external suppliers are assumed to be stochastic, sequential and independent of the system state. External demands follow independent compound Poisson processes. The demand process faced by any internal stage can be determined by the bill of materials, and it is still compound Poisson due to the continuous-time base-stock policy. Demands at each stage are satisfied as much as possible from on-hand inventory. The unsatisfied demands are fully backordered and are satisfied on a FCFS basis as the on-hand inventory becomes available. Each stage converts possibly multiple items into one final item. We assume that one unit of the final item at each stage requires only one unit of each input item. Generalizations of this assumption is discussed in Section 7. Lastly, the service requirement of the external customers at one stage can be specified by a committed service time and a type 2 fill rate.

The supply chain can be mapped into a graph $(\mathcal{N}, \mathcal{A})$ with the node set $\mathcal{N}$ and edge set $\mathcal{A}$. The nodes represent the stages in the supply chain, and are denoted by $1, \cdots, K$. An arc in $\mathcal{A}$ represents a pair of nodes $i, k \in \mathcal{N}$ that have the demand and supply relationship, and are denoted by $(i, k) \in \mathcal{A}$. It is convenient to assign an index n to each unit in a demand faced by any node, so that the smaller the n, the demand unit has higher priority and therefore should be satisfied prior to other units in the same demand but with larger indices. We define the following notation,

- $X_k(n)$: the backorder delay for the nth unit of a demand at node k.
- $W_k(n)$: the inventory holding time of the corresponding item at stage k that satisfies the nth unit of a demand.
- $L_k(n)$: the total replenishment lead-time for the nth unit of an order placed by node k.
- P_k: the processing cycle time at node k.

- $t_{i,k}$: the transportation lead-time from node i to k, $(i,k) \in \mathcal{A}$.
- S_k: the maximum of the lead-times from external suppliers at node k. $S_k = 0$, if node k does not have an external supplier.
- h_k: inventory holding cost per item per unit of time at node k.
- s_k: base-stock level at node k.
- λ_k: demand rate at node k.
- D_k: demand size at node k. D_k is an integer-valued random variable with $Pr\{D_k > 0\} = 1$.

If node k faces external demand, then we define τ_k and β_k to be the committed service time and the type 2 fill rate at node k respectively. Among these parameters, P_k, $t_{i,k}$, S_k, λ_k, D_k, τ_k and β_k are inputs; $X_k(n)$ or s_k are decision variables. Let $Dmax_k$ be the maximum possible value that D_k can take. According to conventions, we denote $a^+ = \max\{a, 0\}$, and let $E(\cdot), V(\cdot)$ and $\sigma(\cdot)$ be the mean, variance, and standard deviation of a random variable, respectively. In the following sections, we characterize the backorder delays for different units in the same demand in various network topologies. For simplicity, we call pure assembly systems by assembly systems and pure distribution system by distribution systems.

3.1 Analysis of A Single Stage

Consider a stage $k \in \mathcal{N}$. Suppose a demand of size y arrives at time t, we ask the following two questions: (1) when is the corresponding order placed at this stage that satisfies the nth unit of this demand? where $1 \leq n \leq y$; (2) what is the index of the unit in the corresponding order that satisfies the nth unit of this demand? In this section, we develop an approach based on the *backward method* of Zhao and Simchi-Levi (2005) to address these questions. While Zhao and Simchi-Levi (2005) considers Poisson demand and therefore only needs to focus on question (1) for $n = 1$, here we need to address both questions for compound Poisson demand.

We first define the following notations with respect to t. Let $D_{k,1}, D_{k,2}, \ldots,$ be the sizes of demand arrivals prior to t, where $D_{k,1}$ is the size of the most recent demand prior to t, $D_{k,2}$ is the size of the second most recent demand, and so on. In a similar vein, let $\nu_{k,1}, \nu_{k,2}, \ldots,$ be the demand interarrival times prior to t, where $\nu_{k,1}$ is the time between the most recent demand and t,

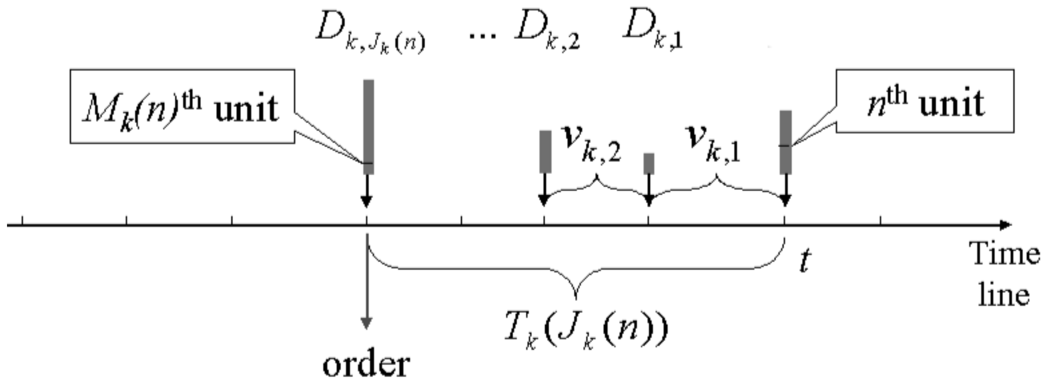

Figure 1: The time line of a single-stage system.

$\nu_{k,2}$ is the time between the second most recent demand and the most recent demand, and so on. See Figure 1 for a visual aid.

Clearly, if $n > s_k$, the corresponding order for the nth unit of the current demand must be placed at time t, i.e., the corresponding order is triggered by the current demand. This is true because the inventory position, s_k, is just enough to satisfy up-to the s_kth unit of the current demand. To answer the second question, we first note that there are s_k items on-hand or incoming right before t. Then the assumption of stochastic sequential lead-times implies that the nth unit of the current demand will be satisfied by the mth item in the corresponding order, where $m = n - s_k$.

If $n \leq s_k$ but $n + D_{k,1} > s_k$, then the corresponding order for the nth unit of the current demand must be placed at time $t - \nu_{k,1}$. To see this, first note that $n \leq s_k$ implies that the inventory position right before t is enough to cover the nth unit, thus the corresponding order must be placed at or before $t - \nu_{k,1}$. On the other hand, $n + D_{k,1} > s_k$ implies that the inventory position right before $t - \nu_{k,1}$ is not sufficient to cover the nth unit, hence the corresponding order must be placed at or after $t - \nu_{k,1}$. To summarize, the corresponding order must be placed exactly at $t - \nu_{k,1}$. To identify the unit in the corresponding order that satisfies the nth unit, we combine $D_{k,1}$ with y into one demand. Then the nth unit in the current demand becomes the $D_{k,1} + n$th unit in the combined demand (due to FCFS). Since there are s_k units on-hand and incoming right before $t - \nu_{k,1}$, the nth unit in the demand y is satisfied by the mth unit in the corresponding order, where $m = D_{k,1} + n - s_k$.

For $j = 2, 3, \ldots, s$, if $n + D_{k,1} + \cdots + D_{k,j-1} \leq s_k$ but $n + D_{k,1} + \cdots + D_{k,j} > s_k$, then the corresponding order for the nth unit of the current demand must be placed at time $t - \nu_{k,1} - \ldots - \nu_{k,j}$. This is true because $n + D_{k,1} + \cdots + D_{k,j-1} \leq s_k$ implies that the inventory position right before

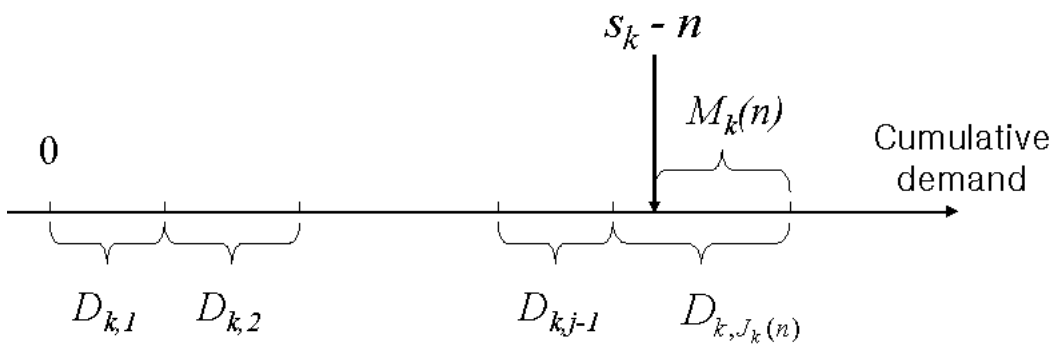

Figure 2: The renewal process generated by $\{D_{k,j}, j \geq 1\}$.

$t-\nu_{k,1}-\cdots-\nu_{k,j-1}$ is enough to cover the nth unit of the current demand, thus the corresponding order must be placed at or before $t-\nu_{k,1}-\cdots-\nu_{k,j}$. On the other hand, $n+D_{k,1}+\cdots+D_{k,j} > s_k$ implies that the inventory position right before $t-\nu_{k,1}-\cdots-\nu_{k,j}$ is not sufficient to cover the nth unit, hence the corresponding order must be placed at or after $t-\nu_{k,1}-\cdots-\nu_{k,j}$. To identify the unit in the corresponding order that satisfies the nth unit, we combine $D_{k,j}, D_{k,j-1}, \ldots, D_{k,1}$ and y into one demand, then the nth unit in the current demand becomes the $D_{k,j}+\ldots+D_{k,1}+n$th unit in the combined demand (due to FCFS). Since there are s_k units on-hand and incoming right before $t-\nu_{k,1}-\ldots-\nu_{k,j}$, the nth unit in demand y is satisfied by the mth unit in the corresponding order, where $m = D_{k,1}+\cdots+D_{k,j}+n-s_k$.

The analysis so far is similar, in principle, to that of Zipkin (1991); except that the later is based on the *forward method*, i.e., identifying the demand unit that will be satisfied by the order triggered by the current demand unit. The analysis introduced here is based on the *backward method*, i.e., for each demand unit, identifying the corresponding unit in the corresponding order that satisfies this demand unit. As we will see, it can be applied to tree supply networks under either continuous-review or periodic-review base-stock policies.

For any $1 \leq n \leq y$, we define random variable $J_k(n)$ so that the corresponding order for the nth unit of the current demand is placed at time $t - T_k(J_k(n))$ (see Figure 1) where

$$T_k(J_k(n)) \;=\; \sum_{j=1}^{J_k(n)} \nu_{k,j}. \tag{1}$$

We also define $M_k(n)$ to be the index of the unit in the corresponding order that satisfies the nth unit.

In view of the above analysis, $J_k(n)$ and $M_k(n)$ can be characterized as follows. Let $\{N_k(i), i \geq$

$0\}$ be the renewal process generated by the demand size process $\{D_{k,j}, j \geq 1\}$ (see Figure 2). Then,

$$J_k(n) = \begin{cases} 0 & \text{if } n > s_k \\ N_k(s_k - n) + 1 & \text{otherwise,} \end{cases} \tag{2}$$

where $J_k(n)$ can choose any value from $\{0, 1, 2, \ldots, s_k - n + 1\}$.

$M_k(n)$ is related to the *remaining life process* $\{O_k(i), i \geq 0\}$ associated with $\{N_k(i), i \geq 0\}$ (see Kulkarni 1995, page 433, for a definition).

$$M_k(n) = \begin{cases} n - s_k, & \text{if } n > s_k \\ O_k(s_k - n), & \text{otherwise.} \end{cases} \tag{3}$$

Given n, both $M_k(n)$ and $J_k(n)$ depend on the base-stock level s_k and the demand size process $\{D_{k,j}, j \geq 1\}$. The joint distribution of $M_k(n)$ and $J_k(n)$ can be characterized by

$$Pr\{M_k(n) = m, J_k(n) = 0\} = 1_{\{n > s_k, m = n - s_k\}} \tag{4}$$

$$Pr\{M_k(n) = m, J_k(n) = j\} = \sum_{l=j-1}^{s_k - n} Pr\{D_{k,1} + \ldots + D_{k,j-1} = l\} Pr\{D_{k,j} = m - n + s_k - l\}, \tag{5}$$

$$m = 1, 2, \ldots, Dmax_k; j = 1, 2, \ldots, s_k - n + 1.$$

In the case of $n \leq s_k$, the dependence between $J_k(n)$ and $M_k(n)$ is determined by the dependence between $N_k(s-n)$ and $O_k(s-n)$. In the special case of Poisson demand, it is easily seen from Eqs. (2)-(3) that $Pr\{M_k(n) = 1\} = 1$ and $Pr\{J_k(n) = s_k\} = 1$ (see also Simchi-Levi and Zhao 2005).

The backorder delay for the nth unit of a demand and the inventory holding time for the corresponding item that satisfies this unit are given by,

$$X_k(n) = [L_k(M_k(n)) - T_k(J_k(n))]^+, \tag{6}$$

$$W_k(n) = [T_k(J_k(n)) - L_k(M_k(n))]^+. \tag{7}$$

Note that Eqs. (6)-(7) depend on the arrival time of the demand. We suppress the dependence without causing confusion. The total replenishment lead-time, $L_k(M_k(n))$, in Eqs. (6)-(7), depends on the backorder delay(s) of the $M_k(n)$th unit of the order placed by node k to its immediate upstream stage(s) at time $t - T_k(J_k(n))$ (see Figure 1). Applying the method of this section to

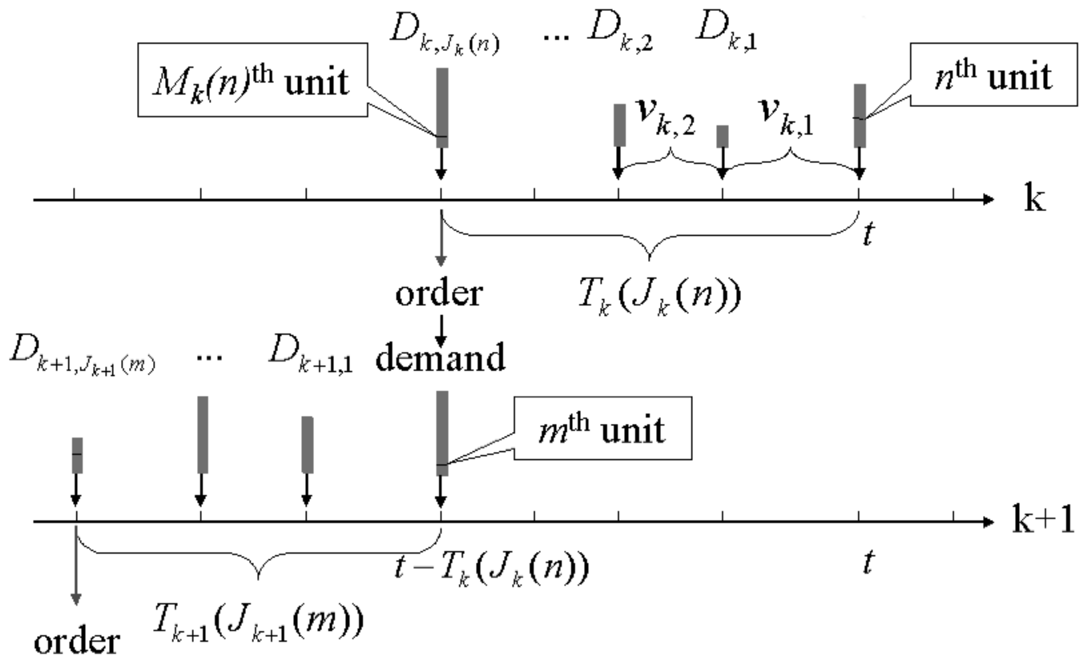

Figure 3: Time line of a serial system.

each of upstream nodes, we can characterize their backorder delays. However, $L_k(\cdot)$ also depends on the network topology, which will be analyzed in the following subsections. The key idea of our approach is that for each unit of an external demand, we identify the corresponding order as well as the corresponding unit in that order placed at each stage of the supply chain that eventually satisfies this unit of demand.

3.2 Serial Systems

Consider a serial supply chain with nodes $k = 1, 2, \cdots, K$ where node K receives external supplies, node k supplies node $k-1$, and node 1 faces external demand. For any node k, consider the nth unit of a demand that arrives at time t. $X_k(n)$ and $W_k(n)$ are determined by Eqs. (6)-(7). Since the order placed by node k is a demand at node $k+1$, $L_k(m)$ in Eqs. (6)-(7) for any $M_k(n) = m$ is given by,

$$L_k(m) = \begin{cases} S_K, & \text{if } k = K \\ X_{k+1}(m) + t_{k+1,k} + P_k, & \text{otherwise,} \end{cases} \tag{8}$$

where $X_{k+1}(m)$ is the backorder delay of the mth unit in the order placed by node k (the demand received by stage $k+1$) at time $t - T_k(J_k(n))$. $X_{k+1}(m)$ can be characterized in the same way as $X_k(n)$. See Figure 3 for a visual aid.

Unlike the supply chains with Poisson demand (see Simchi-Levi and Zhao 2005, Proposition 3.8), $L_k(M_k(n))$ now depends on $T_k(J_k(n))$ because $M_k(n)$ depends on $J_k(n)$. Nevertheless, note that $T_k(\cdot)$, $k = 1, 2, \ldots, K$ are the sums of non-overlapping demand interarrival times (Figure 3), it follows from the compound Poisson demand and the *transit time* assumptions that the serial supply chain with compound Poisson demand can be decomposed into K single stage systems, as follows: we first characterize $X_K(n)$ for all possible $1 \leq n$ by Eqs. (6) and (8); then we determine $L_{K-1}(n)$ for all n by Eq. (8). $X_{K-1}(n)$ can next be characterized by Eq. (6) and the joint distribution of $M_k(n)$ and $J_k(n)$ (see Eqs. 4-5). We can repeat the procedure until $k = 1$.

Chen and Zheng (1994) and Gallego and Zipkin (1999) provide exact characterizations of serial supply chains facing compound Poisson demand. The difference between their approaches and ours is that we focus on the backorder delay of *each unit in a demand* while they focus on the backorders. For phase-type transit times and demand sizes, Zipkin (1991) gives an exact analysis of the probability distributions of the backorder delays in serial and distribution systems.

3.3 Assembly Systems

Consider a pure assembly system where nodes $k = 1, 2, \ldots, K$ supply node 0, and node 0 is the only customer of each node k. Following Song and Zipkin (2002), we make the *committed stock* assumption, i.e., when a demand arrives and some of its required components are in stock but others are not, we put the in-stock components aside as "committed stock".

For node 0, consider the nth unit of a demand that arrives at time t. By Eqs. (6)-(7),

$$X_0(n) = [L_0(M_0(n)) - T_0(J_0(n))]^+, \tag{9}$$

$$W_0(n) = [T_0(J_0(n)) - L_0(M_0(n))]^+, \tag{10}$$

where $L_0(m)$ for any $M_0(n) = m$ is given by

$$L_0(m) \ = \max_{k=1,2,\ldots,K} \{X_k(m) + t_{k,0}\} + P_0, \tag{11}$$

and $X_k(m) = [L_k(M_k(m)) - T_k(J_k(m))]^+$ for all k.

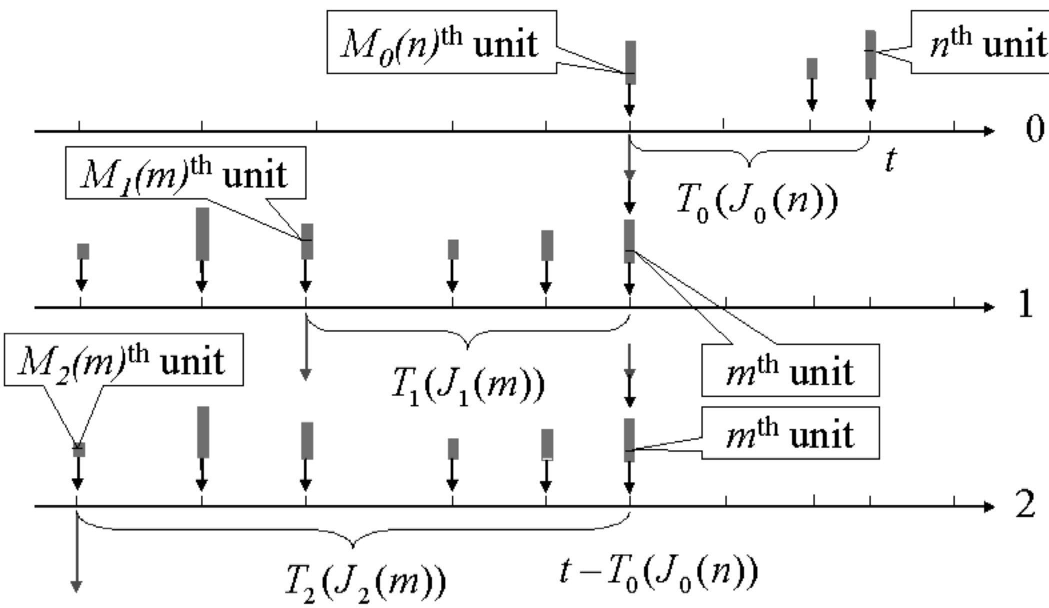

Figure 4: Time line of an assembly system.

First note that all nodes $k = 1, 2, \ldots, K$ receive the corresponding order placed by node 0 (that satisfies the nth unit of the demand at t) at the same time, $t - T_0(J_0(n))$ (see Figure 4). Thus, $T_0(\cdot)$ is not overlapping with $T_k(\cdot)$ for any k. It follows from the compound Poisson demand and the *transit time* assumptions that one can first determine $L_0(m)$ for all m, and then characterize $X_0(n)$ and $W_0(n)$ by Eqs. (9)-(10).

To characterize $L_0(m)$, we need to consider the maximum of $X_k(m) + t_{k,0}$ where $X_k(m)$, $k = 1, 2, \ldots, K$, are dependent because they share the identical index m, and the identical demand process prior to $t - T_0(J_0(n))$ (see Figure 4). To demonstrate the dependence, we define $D_{0,j}$ and $\nu_{0,j}$ for $j \geq 1$ at node 0 in the same way as $D_{k,j}$ and $\nu_{k,j}$ in Section 3.1 but with respect to $t - T_0(J_0(n))$. We also define $\{N_0(i), i \geq 0\}$ to be the renewal process generated by $\{D_{0,j}, j \geq 1\}$, and $\{O_0(i), i \geq 0\}$ to be the associated remaining life process. Lastly, we index the supplying nodes so that $s_1 \leq s_2 \leq \cdots \leq s_K$.

First, $J_k(m)$ and $M_k(m)$ are dependent across $k = 1, 2, \ldots, K$ due to the identical demand size process among nodes $k = 1, 2, \ldots, K$. If $m \leq s_1$, then for $1 \leq j_1 \leq j_2 \leq \cdots \leq j_K$ and any $m_1, m_2, \ldots, m_K$,

$$Pr\{J_1(m) = j_1, M_1(m) = m_1, \ldots, J_K(m) = j_K, M_K(m) = m_K\}$$

$$= Pr\{N_0(s_1 - m) = j_1 - 1, O_0(s_1 - m) = m_1, \ldots, N_0(s_K - m) = j_K - 1, O_0(s_K - m) = m_K\}. \quad (12)$$

Note that all components share the same processes, $N_0(\cdot)$ and $O_0(\cdot)$. For other sequence of $j_1, \ldots, j_K$, the probability in Eq. (12) equals zero. If $s_{k+1} \geq m > s_k$ for some k, then Eq. (12) can be simplified by focusing only on nodes $k+1, \ldots, K$.

Second, given $j_1 \leq j_2 \leq \ldots \leq j_K$, $T_k(j_k)$ are dependent across $k = 1, 2, \ldots, K$ due to the identical demand interarrival times among all $k = 1, 2, \ldots, K$. According to Zhao and Simchi-Levi (2005),

$$\begin{aligned} & Pr\{T_1(j_1) = t_1, T_2(j_2) = t_2, \ldots, T_K(j_K) = t_K\} \\ = & Pr\{\sum_{l=1}^{j_1} \nu_{0,l} = t_1\} Pr\{\sum_{l=j_1+1}^{j_2} \nu_{0,l} = t_2 - t_1\} \cdots Pr\{\sum_{l=j_{K-1}+1}^{j_K} \nu_{0,l} = t_K - t_{K-1}\}. \end{aligned} \quad (13)$$

An assembly system with compound Poisson demand is analytically more challenging than an analogous system with Poisson demand because all the component nodes face not only the common demand interarrival times, but also the common demand size process. Therefore, the first dependence (Eq. 12) is unique for the system with compound Poisson demand, even though the second dependence (Eq. 13) holds for both systems.

3.4 Distribution Systems

Consider a pure distribution system where node 0 is the only supplier of multiple customer nodes $k = 1, 2, \cdots, K$. For node k, consider the nth unit of a demand that arrives at time t. By Eqs. (6)-(7),

$$X_k(n) = [L_k(M_k(n)) - T_k(J_k(n))]^+, \quad (14)$$

$$W_k(n) = [T_k(J_k(n)) - L_k(M_k(n))]^+, \quad (15)$$

where $L_k(m)$ for any $M_k(n) = m$ satisfies,

$$L_k(m) \ = \ X_0(m) + t_{0,k} + P_k, \quad (16)$$

and

$$X_0(m) = [L_0(M_0(m)) - T_0(J_0(m))]^+. \quad (17)$$

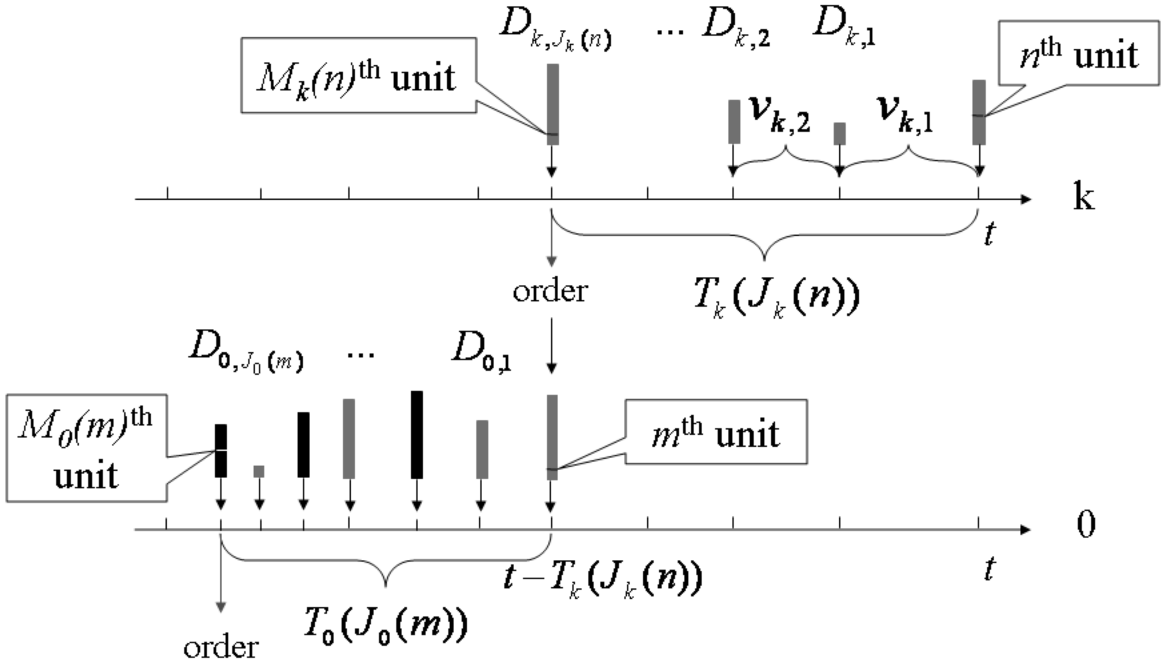

Figure 5: Time line of a distribution system. Bars of different darkness represent demands from different customers

$T_k(\cdot)$, $J_k(\cdot)$ and $M_k(\cdot)$ depend on the demand process at node k, while $T_0(\cdot)$, $J_0(\cdot)$ and $M_0(\cdot)$ depend on the superimposed demand process from all nodes $k = 1, 2, \ldots, K$ (see Figure 5 for an example of $K = 2$). Observe that $T_0(\cdot)$ does not overlap with $T_k(\cdot)$ for any k, the assumptions of independent compound Poisson process and *transit time* implies that one can decompose the distribution system into $K + 1$ single stage systems, as follows: we first characterize $X_0(m)$ and then $L_k(m)$ for all m; then, we determine $X_k(n)$ and $W_k(n)$ for each k and each n.

Zipkin (1991) points out that orders of different node k may experience different stockout delays at node 0. Here we provide a simple characterization. From the above analysis, it is clear that the backorder delays at node 0, $X_0(M_k(n))$, are generally statistically different for the same n but different customer node $k = 1, 2, \ldots, K$, because the distribution of $M_k(n)$ (see Eq. 3), which depends on the demand size process at node k, may vary across different node k. Thus, distribution systems with compound Poisson demand are different from those with Poisson demand, because in the later, the backorder delays at node 0 are statistically the same for all customer nodes k.

3.5 Assembly-Distribution Systems

Consider a set of assembly nodes, if some of their immediate supplying nodes are the same, then we call the network of these assembly nodes and their immediate supplying nodes a assembly-distribution system. Figure 6 depicts two simple examples of such systems. Note that system (a)

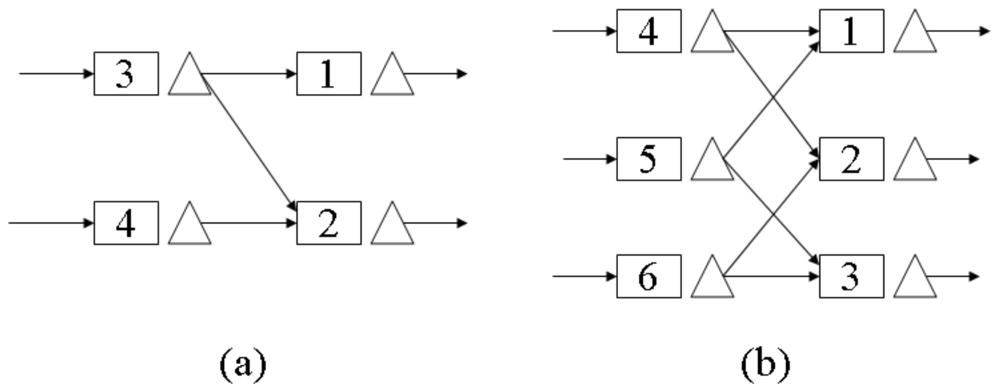

Figure 6: Examples of assembly-distribution systems.

has a tree structure; although system (b) does not have a tree structure, but it satisfies the condition that there is at most one directed path between every two nodes. In this paper, we focus on tree structure supply chains. As we will see, the performance analysis method and approximations, but not the optimization method, can be extended to handle system (b).

In this section, we focus on the systems (a) and (b) in Figure 6. The same method can be applied to more general systems where there is at most one directed path between every two nodes. For node 1 in the system (a), consider the nth unit of a demand that arrives at time t. $X_1(n)$ and $W_1(n)$ can be characterized by Eqs. (6)-(7) if we replace subscript k by 1, where $L_1(m) = X_3(m)+t_{3,1}+P_1$ for any m, and $X_3(m) = [L_3(M_3(m))-T_3(J_3(m))]^+$. As in distribution systems, $T_1(\cdot)$, $J_1(\cdot)$ and $M_1(\cdot)$ depend on the demand process at node 1 while $T_3(\cdot)$, $J_3(\cdot)$ and $M_3(\cdot)$ depend on the superimposed demand process from nodes 1 and 2.

For node 2 in the system (a), consider the nth unit of a demand that arrives at time t. $X_2(n)$ and $W_2(n)$ can be characterized by Eqs. (6)-(7) if we replace subscript k by 2. As in assembly systems, $L_2(m) = \max_{k=3,4}\{X_k(m) + t_{k,2}\} + P_2$ for any m, and

$$X_3(m) \;=\; [L_3(M_3(m)) - T_3(J_3(m))]^+ \tag{18}$$

$$X_4(m) \;=\; [L_4(M_4(m)) - T_4(J_4(m))]^+. \tag{19}$$

However, $J_k(m)$ and $M_k(m)$ (or $T_k(\cdot)$) are dependent across $k = 3, 4$ in a different way from that described by Eq. (12) (Eq. (13), respectively) in pure assembly systems because nodes 3 and 4 face different demand process, i.e., node 3 faces the superimposed demand process from nodes 1 and 2 while node 4 faces the demand process only from node 2. Figure 7 provides a visual aid, where

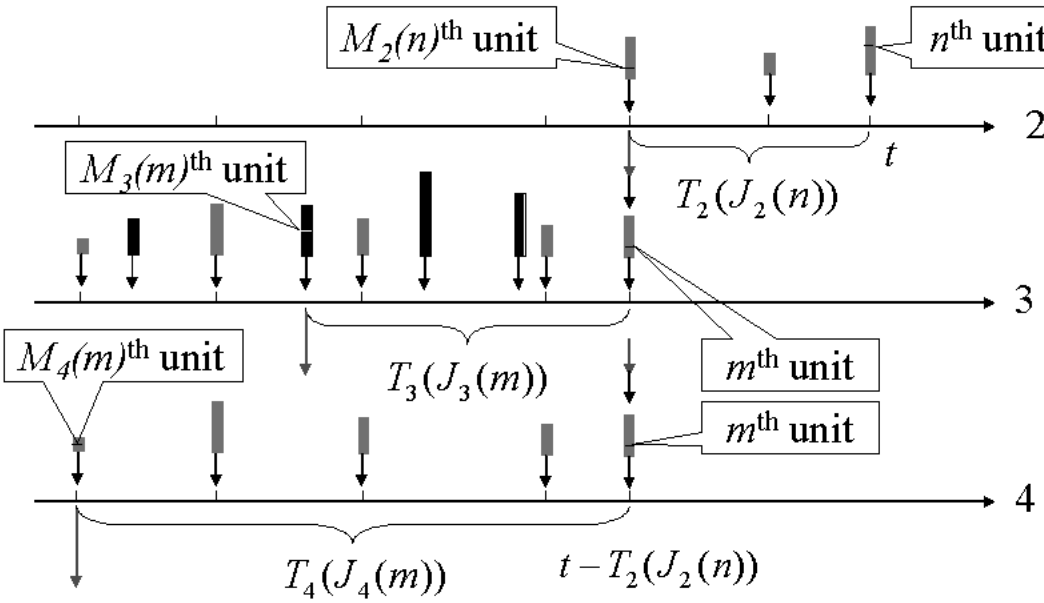

Figure 7: Time line of the assembly-distribution system (a).

the light bars represent demands from node 2, and the dark bars represent demands from node 1. Compare to pure assembly systems (Figure 4), the dependences among the backorder delays of the supplying nodes (e.g., $X_3(\cdot)$ and $X_4(\cdot)$) are weaker in assembly-distribution systems because their demand processes are no longer identical.

The system (b) in Figure 6 can be analyzed in a similar way. For instance, consider node 2 and the nth unit of a demand that arrives at time t. $X_2(n)$ and $W_2(n)$ can be characterized in the same way as in system (a), but now $L_2(m) = \max_{k=4,5,6}\{X_k(m) + t_{k,2}\} + P_2$ for any m, where $X_4(m)$, $X_5(m)$ and $X_6(m)$ depend on the demand process prior to $t - T_2(J_2(n))$ faced by node 4, 5 and 6 respectively. It is important to note that these demand processes are dependent but not identical because they share some but not all common demand arrivals.

3.6 Performance Measures

To determine the cost measures and the service levels at node k, we need to consider an arbitrary demand unit (or a randomly chosen demand unit, equivalently) at this node. Let $p_{k,n}$ be the long-run proportion of demand units that are the nth unit of a demand at node k, it is well known that (e.g., Sigman 2001),

$$p_{k,n} = Pr\{D_k \geq n\}/E(D_k). \tag{20}$$

Define X_k and W_k for an arbitrary demand unit as follows,

$$Pr\{X_k \le x\} = \sum_{n\ge 1} p_{k,n} Pr\{X_k(n) \le x\} \tag{21}$$

$$Pr\{W_k \le w\} = \sum_{n\ge 1} p_{k,n} Pr\{W_k(n) \le w\}. \tag{22}$$

X_k is the backorder delay of an arbitrary demand unit at node k, and W_k is the inventory holding time of the corresponding item at node k that satisfies an arbitrary demand unit. Clearly,

$$E(X_k) = \sum_n p_{k,n} E(X_k(n)), \tag{23}$$

$$E(W_k) = \sum_n p_{k,n} E(W_k(n)). \tag{24}$$

By Little's law, the expected backorders and the expected on-hand inventory at node k are given by

$$E(B_k) = E(X_k)\lambda_k E(D_k), \tag{25}$$

$$E(I_k) = E(W_k)\lambda_k E(D_k). \tag{26}$$

The type 2 fill rate, β_k, within a committed service time τ_k, is given by,

$$\beta_k = \sum_n p_{k,n} Pr\{X_k(n) \le \tau_k\}. \tag{27}$$

If node k is an assembly node, then in addition to the on-hand inventory of the finished good, I_k, we also need to consider the component inventories, I_k^i, $\forall(i,k) \in \mathcal{A}$. These inventories are held at node k without being processed because the corresponding units of other required components are not yet replenished. By Little's law,

$$E(I_k^i) = \lambda_k E(D_k) \sum_n p_{k,n} E(W_k^i(n)), \tag{28}$$

where

$$W_k^i(n) = \max_{\{l|(l,k)\in\mathcal{A}\}} \{X_l(M_k(n)) + t_{l,k}\} - X_i(M_k(n)) - t_{i,k}. \tag{29}$$

The expected value of other types of inventories, i.e., the pipeline inventories due to the processing cycle times and the transportation lead-times are determined by the demand processes and the transit times, and thus they are constants.

4 The Periodic-Review Supply Chains

In this section, we consider the tree structure supply chains where each stage uses a periodic-review base-stock policy. Other assumptions remain the same as in Section 3. The review periods of different stages can be different. Without loss of generality, we assume that the review periods, R_k, $k = 1, 2, \ldots, K$, are nested, i.e., whenever a stage reviews its inventory and makes order decisions, all stages that receives supply from this stage also review inventories and make order decisions (we point out that the same assumption is made by Graves 1985 and Axsater and Rosling 1993 for installation policies).

At each node, demand in one period is satisfied in the order of its arrival while demands in consecutive periods are satisfied on a FCFS basis. We denote demand in a review period at node k to be $\tilde{D}_k$, and denote the maximum possible value that $\tilde{D}_k$ can take to be $\tilde{D}max_k$. By the assumptions that external demands follow independent compound Poisson processes, and the review periods are nested, $\tilde{D}_k$ in different review periods are integer-valued i.i.d random variables, with $Pr\{\tilde{D}_k < 0\} = 0$ and $Pr\{\tilde{D}_k = 0\} > 0$. Clearly, $\tilde{D}_k$ is generated by the compound Poisson process with rate λ_k and size D_k, which is the superposition of the external demand processes faced by the nodes that are either directly or indirectly supplied by node k.

In this paper, we ignore the customer waiting time within one review period, and assume that demands arrive at the end of each review period. This assumption holds true for internal stages when all stages in a supply chain share the same review period. This assumption is also a reasonable approximation when the review periods are sufficiently small with respect to the transportation lead-times and processing cycle times. We refer the reader to Zipkin (2000) Sections 4.3.1 and 9.3.1 for more discussions on modeling demand occuring only at discrete time points. Thus, the demand process at any node k can be viewed as a batch demand process with random size $\tilde{D}_k$ and fixed interarrival time R_k. As in Section 3, each unit in $\tilde{D}_k$ are indexed so that the smaller the index, the higher the priority.

As in the continuous-review supply chains, we can define assembly-distribution systems in periodic-review supply chains. The analysis of such systems is an extension of those of pure assembly and pure distribution systems in Sections 4.3-4.4. Since the extension is similar to that of Section 3.5, we delete the analysis here to save space.

4.1 Analysis of A Single Stage

Consider a single node k. Suppose that the demand realized during period $(t - R_k, t]$ is y, then for the nth unit of this demand $(1 \le n \le y)$, we ask the same two questions as in Section 3.1: (1) when is the corresponding order placed at stage k that satisfies this unit of this demand? (2) what is the index of the unit in the corresponding order that satisfies this unit of demand?

To answer these questions, we define the following notations with respect to t. Let $\tilde{D}_{k,j}$ be the demand realized in the period $(t - (j+1) \times R_k, t - j \times R_k]$ for $j \ge 1$. Note that an order decision is made only at the end of a review period and demand is assumed to arrive only at the end of a review period, thus the only difference between the model in Section 3.1 and the model here is the demand size distribution and the interarrival time process. Applying the analysis of Section 3.1 to demand process with i.i.d. size $\tilde{D}_{k,j}$ and fixed interarrival time R_k yields,

- if $n > s_k$, then the corresponding order is placed at t and the index of the unit in the corresponding order that satisfies this unit of demand is $n - s_k$.
- If $n \le s_k$ and $n + \tilde{D}_{k,1} > s_k$, then the corresponding order is placed at $t - R_k$ and the index of the unit in the corresponding order that satisfies this unit of demand is $\tilde{D}_{k,1} + n - s_k$.
- In general, for $j = 2, 3, \ldots$, if $n + \tilde{D}_{k,1} + \ldots + \tilde{D}_{k,j-1} \le s_k$ and $n + \tilde{D}_{k,1} + \ldots + \tilde{D}_{k,j} > s_k$, then the corresponding order is placed at $t - j \times R_k$ and the index of the unit in the corresponding order that satisfies this unit of demand is $\tilde{D}_{k,1} + \ldots + \tilde{D}_{k,j} + n - s_k$.

For the nth unit of the demand, we define $\tilde{J}_k(n)$ and $\tilde{M}_k(n)$ so that the corresponding order is placed at time $t - \tilde{J}_k(n) \times R_k$, and the index of the unit in the corresponding order that satisfies this unit of demand is $\tilde{M}_k(n)$. Also define the processes $\{\tilde{N}_k(i), i \ge 0\}$ and $\{\tilde{O}_k(i), i \ge 0\}$ to be the renewal process generated by $\{\tilde{D}_{k,j}, j \ge 1\}$, and the associated remaining life process respectively. Then

$$\tilde{J}_k(n) = \begin{cases} 0, & \text{if } n > s_k, \\ \tilde{N}_k(s_k - n) + 1, & \text{otherwise,} \end{cases} \tag{30}$$

and

$$\tilde{M}_k(n) = \begin{cases} n - s_k, & \text{if } n > s_k, \\ \tilde{O}_k(s_k - n), & \text{otherwise.} \end{cases} \tag{31}$$

Unlike the continuous-review systems (Section 3.1), $Pr\{\tilde{D}_{k,j} = 0\} > 0$, thus we do not have an upper bound for $\tilde{J}_k(n)$. Clearly, $\tilde{J}_k(n)$ depends of $\tilde{M}_k(n)$ because $\tilde{N}_k(s_k - n)$ depends on $\tilde{O}_k(s_k - n)$. In view of Eqs. (4)-(5), the joint distribution of $\tilde{J}_k(n)$ and $\tilde{M}_k(n)$ can be characterized as follows:

$$Pr\{\tilde{M}_k(n) = m, \tilde{J}_k(n) = 0\} = 1_{\{n>s_k, m=n-s_k\}} \tag{32}$$

$$Pr\{\tilde{M}_k(n) = m, \tilde{J}_k(n) = j\} = \sum_{l=0}^{s_k-n} Pr\{\tilde{D}_{k,1} + \ldots + \tilde{D}_{k,j-1} = l\} Pr\{\tilde{D}_{k,j} = m - n + s_k - l\}, \tag{33}$$

$$m = 1, 2, \ldots, \tilde{D}max_k; j = 1, 2, \ldots.$$

Define $\tilde{T}_k(\tilde{J}_k(n)) = \tilde{J}_k(n) \times R_k$, and let $\tilde{L}_k(n)$, $\tilde{X}_k(n)$ and $\tilde{W}_k(n)$ be the periodic-review system counter-part of $L_k(n)$, $X_k(n)$ and $W_k(n)$, respectively. Then for the nth unit of the demand, it follows from Eqs. (6)-(7) that

$$\tilde{X}_k(n) = [\tilde{L}_k(\tilde{M}_k(n)) - \tilde{T}_k(\tilde{J}_k(n))]^+, \tag{34}$$

$$\tilde{W}_k(n) = [\tilde{T}_k(\tilde{J}_k(n)) - \tilde{L}_k(\tilde{M}_k(n))]^+. \tag{35}$$

4.2 Serial Systems

Consider the serial supply chain in Section 3.2 but assume that each node uses a periodic-review base-stock policy. By the nested review periods assumption, R_{k+1} is an integer multiple of R_k for any $k \geq 1$. For the nth unit of the demand realized in period $(t - R_k, t]$ at node k, $\tilde{X}_k(n)$ and $\tilde{W}_k(n)$ are given by Eqs. (34)-(35). $\tilde{L}_k(m)$ for any $\tilde{M}_k(n) = m$ is determined by,

$$\tilde{L}_k(m) \ = \ \tilde{X}_{k+1}(\aleph_{k+1,k}(m)) + t_{k+1,k} + P_k, \tag{36}$$

where index $\aleph_{k+1,k}(m)$ is defined as follows: $\aleph_{k+1,k}(m) = m$ if $R_{k+1} = R_k$; otherwise, $\aleph_{k+1,k}(m) = \sum_{j=1}^{r-1} \tilde{D}_{k,j} + m$ where r is uniformly distributed in $\{1, \ldots, R_{k+1}/R_k\}$, and $\{\tilde{D}_{k,j}, j \geq 1\}$ is defined in the same way as that in Section 4.1 but here with respect to $t - \tilde{T}_k(\tilde{J}_k(n))$.

Eq. (36) is different from its continuous-review counter-part, Eq. (8), in $\aleph_{k+1,k}(\cdot)$, due to the review periods. On the other hand, similar to the continuous-review systems, $\tilde{L}_k(\tilde{M}_k(n))$ here depends on $\tilde{T}_k(\tilde{J}_k(n))$ because $\tilde{M}_k(n)$ depends on $\tilde{J}_k(n)$. Finally, we can also decompose the serial supply chain under periodic-review base-stock policies into K single-stage systems by characterizing $\tilde{L}_k(n)$ for all n sequentially from $k = K$ to $k = 1$.

4.3 Assembly Systems

Consider the assembly system in Section 4.3 where nodes $k = 1, 2, \ldots, K$ supply the assembly node 0, and each node uses a periodic-review base-stock policy. By the nested review periods assumption, $R_k, k = 1, 2, \ldots, K$ are integer multiples of R_0. For the nth unit of the demand realized in the review period $(t - R_0, t]$ at node 0, Eqs. (34)-(35) imply that

$$\tilde{X}_0(n) = [\tilde{L}_0(\tilde{M}_0(n)) - \tilde{T}_0(\tilde{J}_0(n))]^+, \tag{37}$$

$$\tilde{W}_0(n) = [\tilde{T}_0(\tilde{J}_0(n)) - \tilde{L}_0(\tilde{M}_0(n))]^+, \tag{38}$$

where $\tilde{L}_0(m)$ for any $\tilde{M}_0(n) = m$ is given by,

$$\tilde{L}_0(m) \;=\; \max_{k=1,2,\ldots,K} \{\tilde{X}_k(\aleph_{k,0}(m)) + t_{k,0}\} + P_0, \tag{39}$$

and $\tilde{X}_k(m')$ for any $\aleph_{k,0}(m) = m'$ is

$$\tilde{X}_k(m') \;=\; [\tilde{L}_k(\tilde{M}_k(m')) - \tilde{T}_k(\tilde{J}_k(m'))]^+. \tag{40}$$

In this case, $\aleph_{k,0}(m) = m$ if $R_k = R_0$; otherwise, $\aleph_{k,0}(m) = \sum_{j=1}^{r_k-1} \tilde{D}_{0,j} + m$ where r_k is uniformly distributed in $\{1, \ldots, R_k/R_0\}$, and $\{\tilde{D}_{0,j}, j \geq 1\}$ is defined in the same way as $\{\tilde{D}_{k,j}, j \geq 1\}$ in Section 4.1 but here with respect to $t - \tilde{T}_0(\tilde{J}_0(n))$. We also define $\{\tilde{N}_0(i), i \geq 0\}$ to be the renewal process generated by $\{\tilde{D}_{0,j}, j \geq 1\}$, and $\{\tilde{O}_0(i), i \geq 0\}$ to be the associated remaining life process. Similar to the continuous-review systems in Section 3.3, one can first determine $\tilde{L}_0(m)$ for all m, and then characterize $\tilde{X}_0(n)$ and $\tilde{W}_0(n)$ by Eqs (37)-(38).

To characterize $\tilde{L}_0(m)$, we note that all nodes $k = 1, 2, \ldots, K$ receive the order placed by node 0 that satisfies the nth unit of the demand realized in $(t - R_0, t]$ at the same time, $t - \tilde{T}_0(\tilde{J}_0(n))$. Consider first the case of $R_k = R_0$ for $k = 1, 2, \ldots, K$, thus $\aleph_{k,0}(m) = m$. Clearly, $\tilde{X}_k(m)$ are dependent for $k = 1, 2, \ldots, K$ because the nodes $k = 1, 2, \ldots, K$ face the identical demand size process, and therefore, $\tilde{J}_k(m)$ and $\tilde{M}_k(m)$ are dependent across k. To characterize the dependence, we index the nodes $k = 1, 2, \ldots, K$ so that $s_1 \leq s_2 \leq \cdots \leq s_K$. If $m \leq s_1$, then for $1 \leq j_1 \leq j_2 \leq \cdots \leq j_K$ and any $m_1, m_2, \ldots, m_K$,

$$\begin{aligned} & Pr\{\tilde{J}_1(m) = j_1, \tilde{M}_1(m) = m_1, \ldots, \tilde{J}_K(m) = j_K, \tilde{M}_K(m) = m_K\} \\ = & Pr\{\tilde{N}_0(s_1 - m) = j_1 - 1, \tilde{O}_0(s_1 - m) = m_1, \ldots, \tilde{N}_0(s_K - m) = j_K - 1, \tilde{O}_0(s_K - m) = m_K\} \end{aligned} \tag{41}$$

For other sequence of $j_1, \ldots, j_K$, the probability in Eq. (41) equals zero. Eq. (41) is similar in principle to Eq. (12) of the continuous-review systems. However, unlike the continuous-review systems, $\tilde{T}_k(j_k), k = 1, \ldots, K$ are not no longer dependent for any given sequence of j_k due to the fixed review periods.

In the case of non-identical R_k for $k = 0, 1, \ldots, K$, the analysis is significantly complicated by many possible combinations of $r_k, k = 1, 2, \ldots, K$, as well as the non-identical demand size process $\{\tilde{D}_{k,j}\}$ at different node k. Exact analysis (not reported here) is feasible but long and tedious. Nevertheless, $\tilde{J}_k(\aleph_{k,0}(m))$ and $\tilde{X}_k(\aleph_{k,0}(m))$ are still dependent across k due to the identical order process $\{\tilde{D}_{0,j}, j \geq 1\}$ placed by node 0 to all nodes $k = 1, \ldots, K$.

4.4 Distribution Systems

Distribution systems under periodic-review policies are different from the continuous-review counterparts because customer nodes can place orders to the distribution node at the same time. Therefore, the distribution node faces the decisions of inventory allocation, i.e., how to allocate its on-hand inventory to the customer nodes when total demand exceeds the supply?

The inventory allocation problem has been studied extensively in the literature, we refer the reader to Jackson (1988), Graves (1996) and reference therein. Graves (1996) introduces the "virtual allocation rule" which works as follows: the distribution node commits inventory to the external demand in the sequence of the demand arrivals rather than the sequence at which the customer nodes place orders. Virtual allocation is not the optimal allocation rule. For independent Poisson external demands, Graves (1996) demonstrates the tractability and effectiveness of this allocation rule, and Axsater (1993b) presents an exact analysis of 2-level distribution systems under this rule. In this section, we consider the distribution system in Section 3.4 but assume that each stage uses an installation periodic-review base-stock policy and the distribution node employs the virtual allocation rule. The objective is to extend the analysis in Sections 4.1-4.3 to these distribution systems facing compound Poisson demand.

Due to the nested review periods assumption, the review period at node 0, R_0, is an integer multiple of R_k for all k. Consider a customer node k, and the nth unit of the demand realized in review period $(t - R_k, t]$. By Eqs. (34)-(35),

$$\tilde{X}_k(n) = [\tilde{L}_k(\tilde{M}_k(n)) - \tilde{T}_k(\tilde{J}_k(n))]^+, \tag{42}$$

$$\tilde{W}_k(n) = [\tilde{T}_k(\tilde{J}_k(n)) - \tilde{L}_k(\tilde{M}_k(n))]^+, \tag{43}$$

where $\tilde{L}_k(m)$ for any $\tilde{M}_k(n) = m$ is given by,

$$\tilde{L}_k(m) \;=\; \tilde{X}_0(\aleph_{0,k}(m)) + t_{0,k} + P_k, \tag{44}$$

and $\tilde{X}_0(m')$ for any $\aleph_{0,k}(m) = m'$ is,

$$\tilde{X}_0(m') \;=\; [\tilde{L}_0(\tilde{M}_0(m')) - \tilde{T}_0(\tilde{J}_0(m'))]^+. \tag{45}$$

$\aleph_{0,k}(m)$ now depends on the review periods as well as the inventory allocation rule. When $R_k = R_0$, $\aleph_{0,k}(m)$ can be determined as follows: let $U_{k,m}$ be the time from the beginning of a review period to the arrival of the mth unit of the demand at node k in the same review period. Furthermore, let $\tilde{D}'_{0,k}(U)$ be the total external demand realized during time U at all customer nodes of node 0 other than node k. According to the virtual allocation rule, $Pr\{\aleph_{0,k}(m) = m + i\} = Pr\{\tilde{D}'_{0,k}(U_{k,m}) = i\}$ for any $i \geq 0$ because $\aleph_{0,k}(m) = m + \tilde{D}'_{0,k}(U_{k,m})$. To characterize $Pr\{\tilde{D}'_{0,k}(U_{k,m}) = i\}$, we first note that $U_{k,m}$ is uniformly distributed in one review period, R_k, by Theorem 5-13 and Corollary 5-13 of Heyman and Sobel (1984). For compound Poisson process, the generating function of the total demand realized in a given time interval is provided by Heyman and Sobel (1984, page 151). Therefore, $Pr\{\tilde{D}'_{0,k}(U_{k,m}) = i\}$ can be computed by conditioning on $U_{k,m}$. When $R_k \neq R_0$, $\aleph_{0,k}(m) = \sum_{j=1}^{r_k-1} \tilde{D}_{k,j} + \tilde{D}'_{0,k}((r_k - 1) \times R_k + U_{k,m})$ where r_k is uniformly distributed in $\{1, \ldots, R_0/R_k\}$, and $\{\tilde{D}_{k,j}, j \geq 1\}$ is defined in the same way as that in Section 4.1 but with respect to $t - \tilde{T}_k(\tilde{J}_k(n))$.

Similar to the continuous-review systems in Section 3.4, $\tilde{M}_k(\cdot)$ and $\aleph_{0,k}(\cdot)$ typically have different distributions for different customer node k. Therefore, the backorder delays at node 0 are statistically different across $k = 1, 2, \ldots, K$. Furthermore, the distribution system here can also be decomposed into $K + 1$ single-stage systems, as follows: we first characterize $\tilde{X}_0(m')$ for each m', which allows us to determine $\tilde{L}_k(m)$ for each m. Then, we compute $\tilde{X}_k(n)$ and $\tilde{W}_k(n)$ by Eqs. (42)-(43).

4.5 Performance Measures

For node k, we define $\tilde{p}_{k,n}$ to be the long-run proportion of demand units that are the nth unit of a demand at this node, where $n \geq 1$. Note that $Pr\{\tilde{D}_k = 0\} > 0$, then

$$\begin{aligned}\tilde{p}_{k,n} &= Pr\{\tilde{D}_k \geq n | \tilde{D}_k > 0\}/E(\tilde{D}_k | \tilde{D}_k > 0), \\ &= Pr\{\tilde{D}_k \geq n\}/E(\tilde{D}_k). \end{aligned} \tag{46}$$

We define $\tilde{X}_k$ and $\tilde{W}_k$ as the counter-part of X_k and W_k (Section 3.6) in the periodic-review systems, that is,

$$Pr\{\tilde{X}_k \leq x\} = \sum_{n \geq 1} \tilde{p}_{k,n} Pr\{\tilde{X}_k(n) \leq x\} \tag{47}$$

$$Pr\{\tilde{W}_k \leq w\} = \sum_{n \geq 1} \tilde{p}_{k,n} Pr\{\tilde{W}_k(n) \leq w\}. \tag{48}$$

Then,

$$E(\tilde{X}_k) = \sum_{n} \tilde{p}_{k,n} E(\tilde{X}_k(n)), \tag{49}$$

$$E(\tilde{W}_k) = \sum_{n} \tilde{p}_{k,n} E(\tilde{W}_k(n)). \tag{50}$$

We can determine the performance measures and service levels at node k, as follows,

$$E(\tilde{B}_k) = E(\tilde{X}_k) \times E(\tilde{D}_k)/R_k, \tag{51}$$

$$E(\tilde{I}_k) = E(\tilde{W}_k) \times E(\tilde{D}_k)/R_k, \tag{52}$$

$$\tilde{\beta}_k = \sum_{n} \tilde{p}_{k,n} Pr\{\tilde{X}_k(n) \leq \tau_k\}. \tag{53}$$

If node k is an assembly node, then the expected component inventories, $E(\tilde{I}_k^i)$, $\forall (i,k) \in \mathcal{A}$ is given by,

$$E(\tilde{I}_k^i) \; = \; \sum_{n} \tilde{p}_{k,n} E(\tilde{W}_k^i(n)) \times E(\tilde{D}_k)/R_k \tag{54}$$

where

$$\tilde{W}_k^i(n) \; = \; \max_{\{l | (l,k) \in \mathcal{A}\}} \{\tilde{X}_l(\aleph_{l,k}(\tilde{M}_k(n))) + t_{l,k}\} - \tilde{X}_i(\aleph_{i,k}(\tilde{M}_k(n))) - t_{i,k}. \tag{55}$$

Finally, $E(\tilde{D}_k)/R_k$ in the above equations can also be expressed by $\lambda_k E(D_k)$.

5 Approximations and Optimization

The analyzes in Sections 3-4 provide bases for exact evaluation and optimization of small size problems. But for larger problem with many stages, exact evaluation is not computationally tractable due to the dependent backorder delays in the assembly and assembly-distribution systems. This is true because the exact characterization of the system performances requires the knowledge of the joint probability distribution of the backorder delays at all stages, and therefore, a tree structure supply chain generally cannot be decomposed into multiple single-stage systems where each stage can characterized separately. On the other hand, the exact analyzes in Sections 3-4 provide bases for developing approximations. Thus, the objective of this section is to develop simple and tractable approximations for both the continuous-review and the periodic-review supply chains, and based on which, we formulate and solve the optimization problems.

5.1 The Continuous-Review Supply Chains

We first ignore the correlations in the assembly and the assembly-distribution systems, i.e., we focus only on the marginal distribution of the backorder delay at each stage of the network rather than the joint distribution. Using this approximation, tree structure supply chains can be decomposed into multiple single-stage systems where each stage can be characterized separately.

Consider node k, given the distribution of $L_k(m)$ for all m, s_k and the demand process λ_k and D_k, we can compute the distribution of $X_k(n)$ for each n by Eq. (6) as follows,

$$Pr\{X_k(n) \le x\} \;=\; \sum_{m\ge 1, j\ge 0} Pr\{M_k(n)=m, J_k(n)=j\} Pr\{L_k(m)-T_k(j)\le x\}. \tag{56}$$

Note that as $s_k - n \to \infty$, $Pr\{M_k(n)=m\} \to p_{k,m}$ (Kulkarni 1995) which does not depend on n. Hence,

$$Pr\{X_k(n)\le x\} \;\to\; \sum_{m\ge 1} p_{k,m} \sum_{j\ge 0} Pr\{J_k(n)=j\} Pr\{L_k(m)-T_k(j)\le x\} \text{ as } s_k\to\infty. \tag{57}$$

Define random variable L_k so that $Pr\{L_k=t\} = \sum_{m\ge 1} p_{k,m} Pr\{L_k(m)=t\}$, then by Eqs. (56)-(57), we can approximate $X_k(n)$ by

$$X_k(n) \;\approx\; (L_k - T_k(J_k(n)))^+, \tag{58}$$

where L_k is independent of $T_k(J_k(n))$ for any n.

In view of the fact that the limiting distribution of $Pr\{M_k(n) = m\}$ is identical to $p_{k,m}$ (the long-run proportion of demand units that are the mth unit of a demand at stage k), L_k is related to the stockout delays of an arbitrary demand unit at the immediate suppliers of node k. Indeed, $p_{k,m} = p_{k',m}$ in systems where node k is the only customer of node k'. In distribution systems where node k' also face demand from other nodes, $p_{k,m} \neq p_{k',m}$ in general. But since computing backorder delay for each demand unit at each stage is time demanding, we use $X_{k'}$ (based on $p_{k',m}$, see Eq. 21) to approximate the backorder delay at node k' experienced by an arbitrary unit in orders of node k (based on $p_{k,m}$). This approximation allows us to focus on X_k rather than $X_k(n)$, and therefore significantly reduce the computing time.

To further improve numerical efficiency, we utilize the following 2-moment approximation (see, e.g., Graves 1985 and Zipkin 1991) for X_k: given $E(L_k)$ and $V(L_k)$, s_k and the demand process λ_k and D_k, we compute $E(X_k)$ and $V(X_k)$ as follows,

1. compute the mean and variance of the lead-time demand, Y_k, where

$$E(Y_k) = E(L_k)\lambda_k E(D_k), \tag{59}$$

$$V(Y_k) = \lambda_k E(D_k^2)E(L_k) + (\lambda_k E(D_k))^2 V(L_k). \tag{60}$$

   Then fit the lead-time demand distribution by a Negative Binomial distribution which matches the first 2 moments.

2. Compute the expected backorders $E(B_k)$ and the variance of the backorders $V(B_k)$ by

$$B_k = (Y_k - s_k)^+. \tag{61}$$

3. Finally, compute $E(X_k)$, $V(X_k)$ and $E(I_k)$ as follows,

$$E(X_k) = E(B_k)/\lambda_k/E(D_k) \tag{62}$$

$$V(X_k) = (V(B_k) - \lambda_k E(D_k^2)E(X_k))/(\lambda_k E(D_k))^2, \tag{63}$$

$$E(I_k) = s_k - E(Y_k) + E(B_k). \tag{64}$$

We refer the reader to Zipkin (1991, 2000) for more discussions.

One important advantage of these approximation is their simplicity, which allows for fast system optimization. However, if stage k is a distribution node, then the backorder delays at node k to all its customers are identical under these approximations. By Eqs. (61)-(62), $E(X_k)$ has a monotonic relationship with s_k. Thus, we can use either $E(X_k)$ or s_k as the decision variable at stage k. Zipkin (1991) studies the accuracy of this 2-moment approximation in serial systems. In Section 6, we will develop insights into the conditions under which the approximations introduced in this section are reasonably accurate in tree structure supply chains.

If node k is supplied by other nodes in the network, say node i, then the replenishment lead-time from node i to node k can be approximated by $X_i + t_{i,k}$. If node k is an assembly node, then by Eqs. (28)-(29), the expected component inventory level at node k can be approximated by

$$E(I_k^i) \;\approx\; \lambda_k E(D_k) E[\max_{l|(l,k)\in\mathcal{A}}\{X_l + t_{l,k}\} - X_i - t_{i,k}], \tag{65}$$

where the mean and variance of the maximum of independent random variables can be computed by Clark's approximation (Clark 1961).

If node k faces external demand, then it follows from Eqs. (27) and (58) that the service level, the type 2 fill rate within τ_k, at node k can be approximated by

$$\sum_{n\geq 1} p_{k,n} Pr\{L_k - T_k(J_k(n)) \leq \tau_k\} \;\geq\; \beta_k. \tag{66}$$

Define $\Delta_n = L_k - \sum_{j=1}^{J_k(n)} \nu_{k,j} - \tau_k$. Given L_k, s_k, τ_k and demand process λ_k and D_k, we approximate Δ_n by a Normal random variable with mean and variance determined as follows (see, e.g., Zipkin (2000) Section C.2.3.8),

$$E(\Delta_n) = E(L_k) - E(J_k(n))/\lambda_k - \tau_k, \tag{67}$$

$$V(\Delta_n) = V(L_k) + E(J_k(n))/\lambda_k^2 + V(J_k(n))/\lambda_k^2. \tag{68}$$

To determine $E(J_k(n))$ and $V(J_k(n))$, we note that if $s_k < n$, then $J_k(n) = 0$ and therefore $E(J_k(n)) = V(J_k(n)) = 0$; if $s_k \geq n$, then by Eq. (2), $J_k(n) = 1 + N_k(s_k - n)$, where $E(N_k(s_k - n))$ and $V(N_k(s_k - n))$ can be approximated by the corresponding asymptotic value $(s_k - n)/E(D_k)$ and $(s_k - n)V(D_k)/E^3(D_k)$ (as $s_k - n \to \infty$) respectively.

We now present the optimization problem for the continuous-review supply chains. To this end, we define $\overline{X}_k$ to be a vector representing the backorder delays at the suppliers of node k, i.e., $\overline{X}_k = \{X_i|(i,k) \in \mathcal{A}\}$. Given the mean and variance of $\overline{X}_k$ ($E(\overline{X}_k)$ and $V(\overline{X}_k)$), $E(X_k)$, and the demand process λ_k and D_k, we can determine s_k, $V(X_k)$ and the safety-stock carrying costs H_k at node k. In particular,

$$H_k(E(\overline{X}_k), V(\overline{X}_k), E(X_k)) \;=\; h_k E(I_k) + \sum_{(i,k)\in\mathcal{A}} h_i E(I_k^i). \tag{69}$$

We can formulate the following mathematic program using $E(X_k), k = 1, 2, \ldots, K$ as decision variables:

$$\begin{aligned}
\mathbf{P} \qquad \min \qquad & \sum_{k=1}^{K} H_k(E(\overline{X}_k), V(\overline{X}_k), E(X_k)) \\
s.t. \qquad & L_k = \max\{X_i + t_{i,k}, \forall (i,k) \in \mathcal{A}\} + P_k, \forall k \in \mathcal{N}, \\
& 0 \le E(X_k) \le \min\{Q_k, E(L_k)\}, \forall k \in \mathcal{N}, \\
& V(X_k) \le \sigma_k^2, \forall k \in \mathcal{N}, \\
& Pr\{X_k \le \tau_k\} \ge \beta_k, \text{for node k serving external customers.}
\end{aligned}$$

The Q_k and σ_k^2 are the maximum allowable expected backorder delay and delay variance at node k, respectively. Program **P** is different from that of Simchi-Levi and Zhao (2005) in two ways: (1) s_k and $V(X_k)$ are computed in a different way by L_k and $E(X_k)$ due to the compound demand processes, and (2) the customer service is computed in a different way. Simchi-Levi and Zhao (2005) develops an algorithm based on dynamic programming to optimally coordinate the continuous-time base-stock policies for supply chains with independent Poisson process and stochastic lead-times. The same dynamic programming algorithm can be applied here. We refer the reader to Zhao and Simchi-Levi (2005) for more discussions on the algorithm.

5.2 The Periodic-Review Supply Chains

As we illustrated in Section 4, the impact of different review periods are complex and problem specific. In the rest of the paper, we focus on an important special case where the review periods of all stages are identical (see also Graves and Willems 2000, 2002). Similar to the continuous-review systems, we ignore the correlations in the assembly and the assembly-distribution systems.

Consider node k and the nth unit of the demand realized in one review period. By Eq. (34),

$$Pr\{\tilde{X}_k(n) \le x\} \;=\; \sum_{m\ge 1, j\ge 0} Pr\{\tilde{M}_k(n) = m, \tilde{J}_k(n) = j\} Pr\{\tilde{L}_k(m) - \tilde{T}_k(j) \le x\}. \qquad (70)$$

Similar to the continuous-review systems in Section 5.1, we observe that as $s_k - n \to \infty$, $Pr\{\tilde{M}_k(n) = m\} \to \tilde{p}_{k,m}$ which does not depend on n. Hence,

$$Pr\{\tilde{X}_k(n) \le x\} \;\to\; \sum_{m\ge 1} \tilde{p}_{k,m} \sum_{j\ge 0} Pr\{\tilde{J}_k(n) = j\} Pr\{\tilde{L}_k(m) - \tilde{T}_k(j) \le x\} \quad \text{as } s_k \to \infty. \qquad (71)$$

Define random variable $\tilde{L}_k$ so that $Pr\{\tilde{L}_k = t\} = \sum_{m\ge 1} \tilde{p}_{k,m} Pr\{\tilde{L}_k(m) = t\}$. Then, by Eqs. (70)-(71), we can approximate $\tilde{X}_k(n)$ by

$$\tilde{X}_k(n) \;\approx\; (\tilde{L}_k - \tilde{T}_k(\tilde{J}_k(n)))^+. \qquad (72)$$

As in the continuous-review systems, we focus on the backorder delay for an arbitrary unit of demand at each node, $\tilde{X}_k$ (Eq. 48). Due to the similarity between Eq. (58) and Eq. (72), we utilize a 2-moment approximation similar to that of Section 5.1. More specifically, given $E(\tilde{L}_k)$ and $V(\tilde{L}_k)$ at stage k, s_k, R_k and the distribution of demand size $\tilde{D}_k$, we compute $E(\tilde{X}_k)$, $V(\tilde{X}_k)$ and $E(\tilde{I}_k)$ as follows,

1. compute the mean and variance of the lead-time demand, $\tilde{Y}_k$, where

$$E(\tilde{Y}_k) = E(\tilde{L}_k)E(\tilde{D}_k)/R_k, \qquad (73)$$

$$V(\tilde{Y}_k) = E(\tilde{L}_k/R_k)V(\tilde{D}_k) + V(\tilde{L}_k/R_k)E^2(\tilde{D}_k) \qquad (74)$$

$$= E(\tilde{L}_k)V(\tilde{D}_k)/R_k + V(\tilde{L}_k)E^2(\tilde{D}_k)/R_k^2. \qquad (75)$$

   Eq. (75) follows from Zipkin (2000) Section C.2.3.8. Then fit the lead-time demand distribution by a Negative Binomial distribution which matches the first 2 moments.

2. Compute the expected backorders $E(\tilde{B}_k)$ and the variance of the backorders $V(\tilde{B}_k)$ by

$$\tilde{B}_k = (\tilde{Y}_k - s_k)^+. \qquad (76)$$

3. Finally,

$$E(\tilde{X}_k) = E(\tilde{B}_k)R_k/E(\tilde{D}_k) \tag{77}$$

$$V(\tilde{X}_k) = (V(\tilde{B}_k)R_k - E(\tilde{X}_k)V(\tilde{D}_k))R_k/E^2(\tilde{D}_k), \tag{78}$$

$$E(\tilde{I}_k) = s_k - E(\tilde{Y}_k) + E(\tilde{B}_k). \tag{79}$$

For compound Poisson demand, Eqs. (73)-(79) reduce to Eqs. (59)-(64). Note that Eqs. (59)-(64) do not depend on R_k. Under the assumptions that all stages have the same review period, and the waiting times within one review period can be ignored, our numerical study in Section 6 indicate that these approximations are sufficiently accurate except when $R_k/E(\tilde{L}_k)$ is relatively large, e.g., > 0.5.

If node k faces external demand, then by Eqs. (53) and (72), we can approximate the type 2 fill rate constraint as follows,

$$\sum_{n \geq 1} \tilde{p}_{k,n} Pr\{\tilde{L}_k - \tilde{T}_k(\tilde{J}_k(n)) \leq \tau_k\} \ \geq \ \beta_k. \tag{80}$$

Let $\tilde{\Delta}_n = \tilde{L}_k - \sum_{j=1}^{\tilde{J}_k(n)} R_k - \tau_k$. We can approximate $\tilde{\Delta}_n$ by a Normal random variable with mean and variance determined as follows,

$$E(\tilde{\Delta}_n) = E(\tilde{L}_k) - E(\tilde{J}_k(n))R_k - \tau_k, \tag{81}$$

$$V(\tilde{\Delta}_n) = V(\tilde{L}_k) + V(\tilde{J}_k(n))R_k^2. \tag{82}$$

To determine $E(\tilde{J}_k(n))$ and $V(\tilde{J}_k(n))$, we note that if $s_k < n$, then $\tilde{J}_k(n) = 0$ and therefore $E(\tilde{J}_k(n)) = V(\tilde{J}_k(n)) = 0$; if $s_k \geq n$, then by Eq. (30), $\tilde{J}_k(n) = 1 + \tilde{N}_k(s_k - n)$, where $E(\tilde{N}_k(s_k - n))$ and $V(\tilde{N}_k(s_k - n))$ can be approximated by the corresponding asymptotic value $(s_k - n)/E(\tilde{D}_k)$ and $(s_k - n)V(\tilde{D}_k)/E^3(\tilde{D}_k)$ (as $s_k - n \to \infty$) respectively.

The inventory costs can be computed in a way similar to those of continuous-review systems, and the optimization program has the same form as program **P**.

6 Numerical Studies

The objective of this section is two-fold: (1) developing insights into the conditions under which the approximations may or may not be sufficiently accurate, and (2) demonstrating the quality of

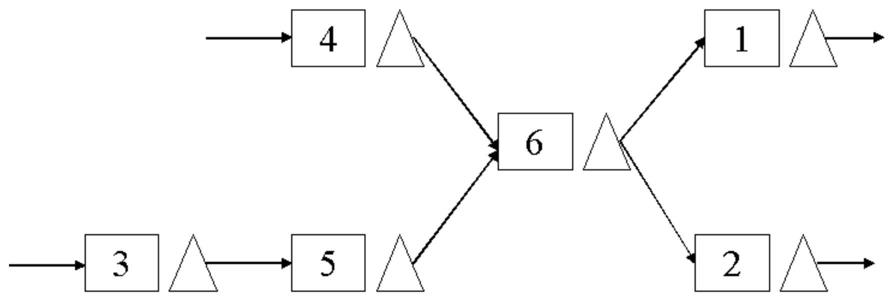

Figure 8: The Example 1.

the solution found by the optimization algorithm. To this end, we consider two numerical examples in the following two subsections.

6.1 A Six-Stage Example

The example is depicted in Figure 8. It consists of one final assembly, node 6, two distribution centers, node 1 and 2, one sub-assembly, node 5, and two component manufacturing sites, node 3 and 4. We assume zero lead-times from external suppliers and zero transportation lead-times between every two nodes, but non-zero and stochastic processing times at all nodes. The processing times follow Erlang distributions with parameters $E(P)$ and n (see, e.g, Zipkin 2000). n here should not be confused with the index of demand units. According to convention, we assume that inventory holding cost increases as one moves downstream the supply chain. In particular, let $h_3 = 1$, $h_4 = 1.5$, $h_5 = 2$, $h_6 = 3$ and $h_1 = h_2 = 4$. Different distribution centers can face different demand rates, demand size distributions and provide different service levels to external customers. Without loss of generality, let $\lambda_1 = 0.7$ and $\lambda_2 = 0.3$. The demand sizes at node 1 and 2 follow discrete Normal distributions with probability density functions $(Pr\{D_1 = 1\}, Pr\{D_1 = 2\}, \cdots, Pr\{D_1 = 7\}) = (0.0062097, 0.0606, 0.2417, 0.383, 0.2417, 0.0606, 0.0062097)$ and $(Pr\{D_1 = 3\}, Pr\{D_1 = 4\}, \cdots, Pr\{D_1 = 9\}) = (0.0062097, 0.0606, 0.2417, 0.383, 0.2417, 0.0606, 0.0062097)$.

To test the accuracy of the approximations, we first use the optimization (DP) algorithm to find a solution and then use a Monte Carlo simulation to evaluate the solution. The Monte Carlo simulation is based on the exact analyzes in Sections 3 and 4. In the simulation, we run 10^4 independent replications for each parameter set and calculate the 95% confidence interval for the performance measures. The confidence intervals for all simulated fill rates are no larger than 1%. The confidence intervals for all the simulated costs are no larger than 3% of their corresponding simulated costs.

For continuous-review systems, we conduct two numerical studies. In the first study, we fix the lead-time parameters $\{n^1, n^2, n^3, n^4, n^5, n^6\} = \{5, 7, 6, 7, 8, 5\}$ and $\{E(P_1), E(P_2), E(P_3), E(P_4), E(P_5), E(P_6)\} = \{6, 6, 1, 2, 5, 6\}$, but vary τ_1, τ_2 and β_1, β_2, where $\tau_k, k = 1, 2$ can choose each value of $\{0, 2, 4\}$, and $\beta_k, k = 1, 2$ can choose each value of $\{0.85, 0.9, 0.95, 0.99\}$. Thus, this study has totally 144 instances. In the second study, we fix $\tau_1 = 1$ and $\tau_2 = 2$. For each $\beta_1 = \beta_2 = 0.85, 0.90, 0.95, 0.99$, we study 100 instances with randomly generated lead-time parameters $\{E(P_1), E(P_2), E(P_3), E(P_4), E(P_5), E(P_6)\}$ and $\{n^1, n^2, n^3, n^4, n^5, n^6\}$, where $E(P_k) \sim Uniform(1, 2, \ldots, 10)$ and $n^k \sim Uniform(1, 2, \ldots, 10)$ for $k = 1, 2, \ldots, 6$.

For each parameter set, we use the DP algorithm to determine the base-stock level at each node and then use simulation to estimate the performance of this solution. Table 1 demonstrates the average and maximum absolute % error in cost, the average and maximum absolute error in β_1 and β_2 for the first study. The absolute percentage error in cost is defined as the absolute difference between the simulated cost and the cost generated by the DP algorithm divided by the simulated cost.

Table 1: The Accuracy of the approximations. Study 1 for the continuous-review systems

Avr. Abs. % Error in Cost	Max. Abs. % Error in Cost	Avr. Abs. Error in β_1	Max. Abs. Error in β_1	Avr. Abs. Error in β_2	Max. Abs. Error in β_2
0.52%	1.7%	0.46%	1.66%	1.24%	3.9%

Table 2 summaries the results of the second study for the continuous-review systems. Tables 1 and 2 demonstrate that the approximations are in general sufficiently accurate for the continuous-review systems under all combinations of the parameters examined in this paper. The largest error is a 3.9% error on β_2 (Table 1), the corresponding instance has $\beta_2 = 0.85$, which indicate that the errors may increase as β decreases. This observation is partially confirmed by Table 2, which shows that the average and maximum errors of the cost and β_2 are increasing as β_1 and β_2 decrease. However, the trend is not clear for the average and maximum errors of β_1.

For periodic-review systems, we conduct two numerical studies identical to those of the continuous-review systems except that in the first study, we also vary the review period $R = 0.5, 1, 2, 4$, and in the second study, we fix $R = 1$ for all stages. The results are summarized in Tables 3 and 4.

Table 2: The Accuracy of the approximations. Study 2 for the continuous-review systems

$\beta_1 = \beta_2 =$	Avr. Abs. % Error in Cost	Max. Abs. % Error in Cost	Avr. Abs. Error in β_1	Max. Abs. Error in β_1	Avr. Abs. Error in β_2	Max. Abs. Error in β_2
0.85	0.67%	3.51%	0.93%	2.48%	2.02%	3.71%
0.90	0.64%	2.54%	0.93%	3.09%	1%	2.53%
0.95	0.46%	1.79%	1.1%	3.28%	0.64%	2.42%
0.99	0.32%	1.36%	0.79%	3.04%	0.45%	1.1%

Table 3: The Accuracy of the approximations. Study 1 for the periodic-review systems

R	Avr. Abs. % Error in Cost	Max. Abs. % Error in Cost	Avr. Abs. Error in β_1	Max. Abs. Error in β_1	Avr. Abs. Error in β_2	Max. Abs. Error in β_2
0.5	0.31%	1.03%	1.48%	3.65%	3.09%	6.44%
1	0.36%	1.46%	1.3%	3.03%	2.76%	5.47%
2	1.32%	3.35%	0.45%	2.32%	2.44%	5.05%
4	4.09%	8.69%	1.49%	3.14%	1.5%	3.11%

Tables 3 and 4 illustrate the average and maximum absolute % error in cost, and the average and maximum absolute error in fill rates for all the combinations of the parameters. Specifically, the average and maximum errors in cost are very small in all cases except when R becomes substantial relative to the processing times, e.g., $R = 4$. At large R, the approximations may perform poorly on cost but still very well on the fill rates. The average errors on type 2 fill rates are relatively small, e.g., $\leq 3.96\%$ (Table 4). However, the maximum errors can be sizeable especially at smaller R, e.g., as large as 6.44% at $R = 0.5$ (Table 3). Interestingly, at R $(= 0.5)$, the errors on cost are the smallest. The accuracies on cost and fill rates are different because we use different approximations for cost (Eq. 64) and fill rate (Eq. 80). Indeed, our numerical study of a single-stage system (not reported here) shows that the 2-moment approximations on expected backorder delays and cost, Eqs. (59)-(64), may subject to large % error when R is large relative to the lead-times, e.g., $R/E(L) > 0.5$; on the other hand, the fill rate approximation, Eq. (80), may perform poorly at small R, e.g., $R/E(L) < 0.1$.

Examination of the instances corresponding to the largest errors in fill-rates (in Table 3) reveals

Table 4: The Accuracy of the approximations. Study 2 for the periodic-review systems

Target $\beta_1 = \beta_2 =$	Avr. Abs. % Error in Cost	Max. Abs. % Error in Cost	Avr. Abs. Error in β_1	Max. Abs. Error in β_1	Avr. Abs. Error in β_2	Max. Abs. Error in β_2
0.85	0.78%	3.48%	2.03%	5.59%	3.96%	7.08%
0.90	0.5%	2.3%	1.11%	3.79%	2.89%	5.98%
0.95	0.35%	1.13%	0.62%	2.73%	1.55%	3.27%
0.99	0.27%	0.89%	0.58%	2.4%	0.11%	0.39%

that the instances also have relatively small target fill-rates (e.g., 85%). This observation is further confirmed by Table 4, which clearly demonstrates that as the target type 2 fill-rates increase from 85% to 99%, both the average and the maximum errors in fill-rates and cost decrease. A careful examination of the instance corresponding to the largest errors in fill-rates in Table 4 also shows that the instances have relatively small n^k (e.g., 1). Note that the coefficient of variation (c.v.) of the processing time P_k equals to $1/\sqrt{n^k}$. The impact of n^k can be explained as follows: the Normal distribution is a poor fit for the random variables representing lead-times when n is close to 1. Indeed, when $n = 1$, the lead-time distribution is exponential, and Clark's method (to compute the maximum of random variables) based on Normal approximation can be far from accurate (Clark 1961).

Given that we ignore the dependency among the lead-times, and we only consider the first two moments of the lead-times, the approximations are reasonably accurate for a range of parameters of interest for both the continuous-review and the periodic-review systems. In particular, the approximations perform well in the important cases when the target type 2 fill-rates are relatively high (e.g., $\geq 95\%$), when the review period R is neither substantial nor negligible comparing to the replenishment lead-time at each stage, and when the lead-time c.v.s are moderate (e.g. < 0.5). One important advantage of the approximations is that they are simple, and therefore they are computationally tractable even for larger size problems (see a 22 node and 21 arcs example in Section 6.2).

We now demonstrate the quality of the solutions obtained by the DP algorithm. To this end, we compare the DP solution to a solution found by a simulation-based search algorithm. Since the simulated fill rates of the DP solution may not closely match with the targets, we adjust the input fill rates and run the DP algorithm repetitively until the target fill rates fall into the 95%

Table 5: Comparison between the solutions found by the DP and the search algorithm. Continuous-review systems

(β_1, β_2)	$\{s_1, s_2, s_3, s_4, s_5, s_6\}$ The search solu.	$\{s_1, s_2, s_3, s_4, s_5, s_6\}$ The DP solu.	Cost (Search)	Cost (DP)	% diff. in costs
$(0.90, 0.90)$	$\{24, 17, 23, 7, 6, 21\}$	$\{27, 20, 6, 0, 15, 21\}$	133.66	140.77	5.32%
$(0.90, 0.96)$	$\{23, 22, 27, 0, 0, 28\}$	$\{22, 23, 12, 11, 18, 21\}$	153.68	161.43	5.04%
$(0.95, 0.90)$	$\{30, 17, 13, 15, 19, 14\}$	$\{33, 20, 6, 0, 15, 21\}$	156.33	163.6	4.65%
$(0.96, 0.96)$	$\{29, 22, 27, 15, 12, 14\}$	$\{28, 22, 6, 0, 15, 31\}$	175.9	182.9	3.98%
$(0.98, 0.95)$	$\{37, 21, 18, 7, 12, 21\}$	$\{36, 21, 6, 0, 9, 34\}$	203.4	208.9	2.7%
$(0.95, 0.99)$	$\{29, 34, 23, 0, 0, 28\}$	$\{28, 33, 6, 0, 9, 34\}$	220.6	222.8	1%
$(0.98, 0.99)$	$\{36, 33, 27, 3, 0, 28\}$	$\{36, 33, 6, 0, 9, 34\}$	248.03	255.47	3%

confidence interval of the simulated fill rates. In the search algorithm, we first identify an upper bound s'_k for the base-stock level at each node. Then, for any base-stock level vector $\{s_3, s_4, s_5, s_6\} \in \bigotimes_{k=3,4,5,6}\{0, \lfloor s'_k/10 \rfloor, \lfloor 2s'_k/10 \rfloor, \cdots, s'_k\}$, we use simulation to evaluate the average system cost and to choose s_1 and s_2 so that the simulated fill-rates closely matches with the targets. Table 5 (Table 6) summarizes the results for the continuous-review systems (the periodic-review systems, respectively) where $\tau_1 = 1, \tau_2 = 2$, $\{E(P_1), E(P_2), E(P_3), E(P_4), E(P_5), E(P_6))\} = \{3, 3, 4, 4, 3, 1\}$ and $\{n^1, n^2, n^3, n^4, n^5, n^6\} = \{6, 5, 2, 8, 6, 3\}$. β_1 and β_2 are allowed to vary. $R = 1$ for the periodic-review systems. All costs are evaluated by simulation.

Table 6: Comparison between the solutions found by the DP and the search algorithm. Periodic-review systems

(β_1, β_2)	$\{s_1, s_2, s_3, s_4, s_5, s_6\}$ The search solu.	$\{s_1, s_2, s_3, s_4, s_5, s_6\}$ The DP solu.	Cost (Search)	Cost (DP)	% diff. in costs
$(0.91, 0.92)$	$\{22, 17, 27, 0, 0, 28\}$	$\{27, 21, 6, 0, 15, 21\}$	136.31	145.08	6.43%
$(0.90, 0.95)$	$\{23, 23, 23, 0, 0, 28\}$	$\{27, 26, 6, 0, 9, 24\}$	155.3	162.4	4.57%
$(0.95, 0.91)$	$\{28, 17, 23, 3, 0, 28\}$	$\{32, 20, 6, 0, 15, 21\}$	155.67	159.35	2.36%
$(0.96, 0.96)$	$\{30, 23, 23, 0, 0, 28\}$	$\{34, 26, 6, 0, 15, 21\}$	181.21	189.1	4.35%
$(0.98, 0.95)$	$\{34, 21, 18, 11, 6, 28\}$	$\{37, 22, 12, 11, 18, 21\}$	209.21	215.84	3.17%
$(0.96, 0.99)$	$\{30, 43, 13, 11, 12, 21\}$	$\{29, 39, 12, 11, 18, 21\}$	261.4	251.35	-3.84%
$(0.98, 0.98)$	$\{36, 29, 23, 3, 0, 28\}$	$\{37, 29, 12, 11, 18, 21\}$	235.2	244.4	3.91%

The percentage difference in costs in Tables 5-6 is defined as the difference between the cost of the DP solution and the cost of the search-based solution divided by the cost of the search-based solution. On a Pentium 1.67 GHZ laptop, the search algorithm takes about 3 hours to solve for one instance while the DP algorithm takes about 2-3 seconds. We first note that in all cases, the cost of the DP solution is reasonably close to that of the search-based solution. The DP solutions tend to perform better as the target fill-rates increase in both the continuous-review and the periodic-review systems. This observation is consistent to our earlier observations on the accuracy of the approximations. Since the approximations are more accurate for higher target fill-rates, the DP algorithm based on the approximations tends to find better solutions. When the fill-rates are in the lower 90%s, the DP solutions can be inferior to the Search solutions by as much as 6.43%. Interestingly, we find that in one case, the DP solution out-performs the search solution. This is possible because the search algorithm only evaluates a subset of the possible base-stock levels.

6.2 A 22 Nodes 21 Arcs Example

In this section, we consider a more elaborate example with 22 nodes and 21 arcs, see Figure 9 for the network structure of the example. This example is inspired by a real world problem, the Bulldozer supply chain (see Graves and Willems 2002). The objective of this section is to further demonstrate the accuracy of the approximations and to show the efficiency of the optimization algorithm.

We use the same inventory holding costs and the expected processing times as those in Graves and Willems. But we consider stochastic processing times, and stochastic, non-zero transportation lead-times with the means generated randomly according to $Uniform\{1, 10\}$, see Appendix III of Simchi-Levi and Zhao (2005) for the costs and lead-times data of the example. The external supply lead-times are zero. We assume that the external demand follows compound Poisson process with $\lambda = 1$, and demand size distribution $(Pr\{D_1 = 1\}, Pr\{D_1 = 2\}, \cdots, Pr\{D_1 = 7\})$ $= (0.0062097, 0.0606, 0.2417, 0.383,\ 0.2417, 0.0606, 0.0062097)$. As in the real world problem, the review period $R = 1$ for all nodes. The target customer service at the final assembly is specified by τ and β where $\tau = 0$.

We study the impact of the lead-time uncertainty and the target β on the accuracy of the approximations. To this end, we assume that all processing times and transportation lead-times follow Erlang distributions with the same coefficient of variation, i.e., the same n. n varies from

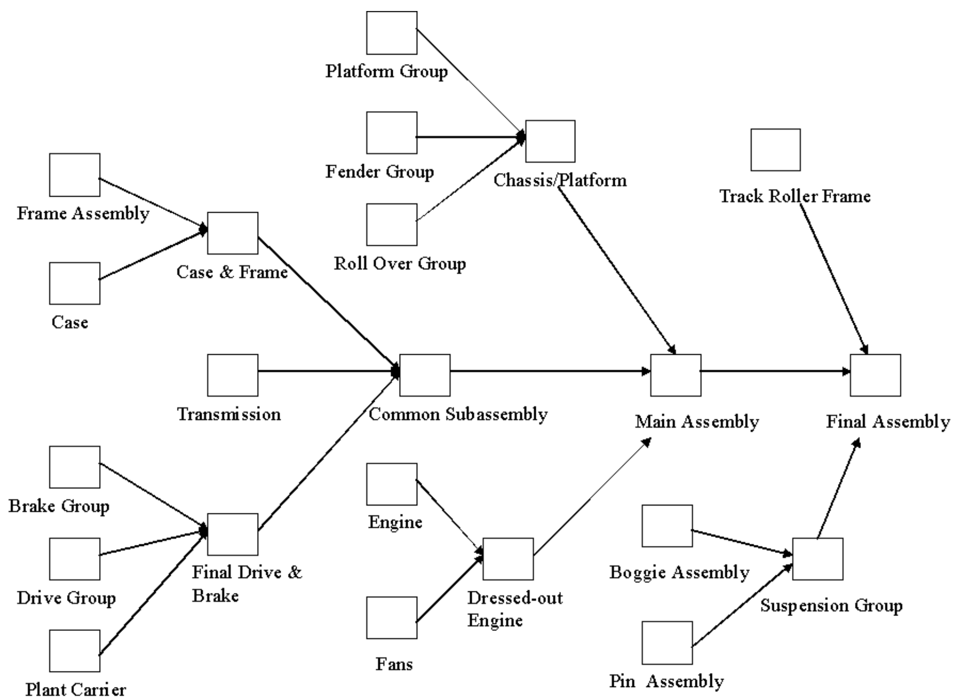

Figure 9: The Example 2.

4, 9 to 16 while $\beta = 0.85, 0.9, 0.95, 0.99$. While this example is computationally prohibitive for the simulation-based search algorithm, it takes around 1 minute for the DP algorithm to generate a solution on a Pentium 1.67 GHZ laptop.

For each parameter set, we use the DP algorithm to determine the base-stock level at each node and then use simulation to estimate the total cost and fill rate. Table 7 presents the absolute percentage difference in costs between simulation and the approximation, and Table 8 shows the absolute difference in fill rates between simulation and the target. The confidence intervals for all simulated fill rates are no larger than 1%, and the confidence intervals for all the simulated costs are no larger than 2% of their corresponding simulated costs.

Table 7: The accuracy of the approximations in cost. Example 2

	$\beta = 0.85\%$	0.90%	0.95%	0.99%
$n = 4$ $(c.v. = 0.5)$	1.99%	0.52%	0.73%	0.18%
9 (0.33)	0.77%	0.94%	0.25%	0.23%
16 (0.25)	0.51%	1.21%	0.13%	0.13%

Table 8: The accuracy of the approximations in fill rate. Example 2

	$\beta = 0.85\%$	0.90%	0.95%	0.99%
$n = 4$ ($c.v. = 0.5$)	5.73%	6.53%	5.58%	3.26%
9 (0.33)	2.28%	3.52%	2.7%	1.6%
16 (0.25)	1.46%	1.65%	1.83%	1.07

Table 7 shows that the cost approximations are sufficiently accurate for all combinations of the lead-time c.v.s and the target fill rates. The approximations become more accurate as the lead-time c.v.s decrease or as the target fill rate increases. Table 8 illustrates that the fill rate approximations are reasonably accurate when the lead-time c.v.s are relatively small or when the target fill rate is high. When the lead-time c.v.s. are relatively large, e.g., > 0.33, the fill rate approximation may perform quite poorly. As we explain before in Section 6.1 that the Clark's method may yield sizeable errors when n is close to 1. Since this example has as many as 8 assembly operations, Clark's method is used repetitively which accumulates errors. Therefore, n has a substantial impact on the accuracy of the approximations in this example.

Given a target service level, e.g., $\tau = 0$ and $\beta = 0.95\%$, we can adjust the input fill rate repetitively for the DP algorithm until the simulated fill rate of the DP solution closely matches the target level. Numerical studies with various lead-time c.v.s reveal some interesting insights. As in Simchi-Levi and Zhao (2005), we first observe that the system-wide inventory cost increases significantly as the lead-time c.v. increases from 0.25, 0.33 to 0.5 (see Table 9). We also observe that the portion of the cost due to component inventories (i.e., I_k^i) is quite substantial, and it tends to increase as the lead-time c.v. increases. The base-stock levels obtained by the DP algorithm is listed in Table 10 in Appendix I.

Table 9: The impact of lead-time uncertainty

n	Lead-time c.v.	Total Cost (Simul.)	Component Inventory Cost/Total Cost
4	0.5	\$2,804,638	29.56%
9	0.33	\$2,046,915	28.46%
16	0.25	\$1,751,459	26.8%

7 Concluding Remarks

In this paper we present an exact and systematic approach to analyze tree-structure supply chains, where the external demands follow independent compound Poisson processes, and the lead-times are stochastic, sequential and exogenously determined. For these supply chains under either continuous-review or periodic-review base-stock policies, we characterize the backorder delay for each unit of a demand at all stages. We analyze the similarity and the structural differences in various material flow topologies between the continuous-review supply chains with compound Poisson demands and those with Poisson demand. We also compare and contrast the continuous-review supply chains with the periodic-review supply chains. In order to efficiently coordinate the inventory policies in the supply chains, we present simple and numerically tractable approximations. Numerical study shows that the approximations are reasonably accurate for a wide range of parameters, and the optimization algorithm can find the optimal or close to optimal solutions efficiently.

Many challenges remain. We conclude the paper by identifying some direct extensions as well as its limitations.

- First, multi-product supply networks with common components and subassemblies, e.g., assemble-to-order systems, are in general difficult to evaluate and coordinate. However, it is interesting to note that these networks can be mapped to single-product supply networks where each stage manages only one product. Here is how: first map each *unique combination of product and facility* into one stage of a single-product supply network, then, set the parameters of the single-product network as follows: the production cycle time of a stage is the production cycle time of the product at the facility in the original system; and the transportation lead-time between two stages is the transportation lead-time between the two corresponding facilties in the orginal system. While some of these single-product systems satisfy the assumptions of this paper, and therefore, can be solved (at least approximately) accordingly, two challenges remain: (1) the mapped single-product system may not have a tree structure, e.g., system (b) of Figure 6; (2) the assumption of identical flow unit may not hold, that is, one unit of different final items at one facility may require different number of units of a common input item.

  It is perhaps worth pointing out that resolving the second challenge requires an extension of the "transit time" model. To see this, let's consider a product at a facility, namely, stage

k in the mapped single-product system, where assembling one unit of the product requires multiple units of a component. Consider the nth unit of a demand that arrives at time t. Since this unit of demand requires multiple units of the component, the corresponding orders of these units of the component may be placed at different times, and therefore, replenished at different times. Note that early replenished units have to wait for the last unit before they are assembled into one unit of the product, which results in additional inventory holding costs. Clearly, exact characterization of these inventory costs requires information about the differences among the lead-times (including the "transit times") of consecutive orders placed by stage k. Therefore, one needs to know the joint distribution of the "transit times" rather than just their marginal distributions.

- Second, we ignore the production and transportation capacity constraints in this paper. Incorporating these constraints into tree structure supply networks pose a substantial challenge. We refer the reader to Parker and Kapuscinski (2004) and Janakiraman and Muckstadt (2004) for the optimal inventory policies in capacitated serial systems, to Glasserman and Tayur (1995) and Kapuscinski and Tayur (1999) for simulation based optimization algorithms, and to Lee and Zipkin (1992), Buzacott, Price and Shanthikumar (1993), Glasserman and Tayur (1996) and Liu, Liu and Yao (2004) for approximations of serial supply chains.

- Thirdly, extending this approach to supply chains where each stage utilizes a batch ordering policy is an important challenge. The difficulty comes from the non-renewal demand processes faced by internal stages if they are the supperpositions of the order processes placed by the downstream stages.

- Finally, more accurate and robust approximations need to be developed for supply chains with long review periods and low target type 2 fill-rates. It is also a challenge to develop accurate approximations for periodic-review supply chains with long and different review periods where demand cannot be assumed to arrive only at the end of each period.

References

Axsater, S. (1993a). Continuous review policies for multi-level inventory systems with stochastic demand. S. Graves, A. Rinnooy Kan, P. Zipkin, eds. *Logistics of Production and Inventory.* Elsevier (North-Holland), Amsterdam. The Netherlands.

Axsater, S. (1993b). Optimization of order-up-to-S policies in two-echelon inventory systems with periodic review. *Naval Research Logistics*, **40**, 245-253.

Axsater, S. (2000). Exact analysis of continuous review (R, Q) policies in two-echelon inventory systems with compound Poisson demand. *Operations Research*, **48**, 686-696.

Axsater, S. (2002). Supply chain operations: serial and distribution inventory systems. Chapter 10 in *Handbooks in Operations Research and Management Science*, Vol. 11: Supply Chain Management: Design, Coordination and Operation. A. G. de Kok and S. C. Graves, eds. North-Holland Publishing Company, Amsterdam, The Netherlands.

Axsater, S. & K. Rosling (1993). Notes: Installation vs. Echelon stock policies for multilevel inventory control. *Management Science*, **39**, 1274-1280.

Buzacott, J., S. Price and J. Shanthikumar (1993). Service level in multi-stage MRP and base-stock controlled production systems. In *New Directions for Operations Research in Manufacturing* (G. Fandel, T. Gulledge, A. Hones, eds.), 445-463. Springer, Berlin, Germany.

Chen, F. (2000). Optimal policies for multi-echelon inventory problems with batch ordering. *Operations Research*, **48**, 376-389.

Chen, F. & Y. S. Zheng (1994). Lower bounds for multi-echelon stochastic inventory systems. *Management Science*, **40**, 1426-1443.

Chen, F. & Y. S. Zheng (1997). One-warehouse multi-retailer systems with centralized stock information. *Operations Research*, **45**, 275-287.

Clark, C. E. (1961). The greatest of a finite set of random variables. *Operations Research*, **9**, 145-162.

Clark, A. J. & H. E. Scarf (1960). Optimal policies for a multi-echelon inventory problem. *Management Science*, **6**, 475-490.

Deuermeyer, B. & L. B. Schwarz (1981). A model for the analysis of system service level in warehouse/retailer distribution systems: the identical retailer case. In *Multilevel production/inventory control systems*, Chap. 13 (TIMS Studies in Management Science 16), L. Schwarz (ed.). Elsevier, New York.

Ettl, M., G. E. Feigin, G. Y. Lin & D. D. Yao (2000). A supply network model with base-stock control and service requirements. *Operations Research*, **48**, 216-232.

Federgruen, A. (1993). Centralized planning models for multi-echelon inventory systems under uncertainty, in: S. .C. Graves et al. (eds.). *Handbooks in OR & MS Vol. 4*, North Holland

Amsterdam, 133-173.

Feigin, G. E. (1999). Inventory planning in large assembly supply chains. In S. Tayur, R. Ganeshan, M. J. Magazine (eds.). *Quantitative Models for Supply Chain Management*, Kluwer Academic Publishers.

Gallego, G. & P. Zipkin (1999). Stock positioning and performance estimation in serial production-transportation systems. *Manufacturing and Service Operations Management*, **1**, 77-88.

Glasserman, P. & S. Tayur (1995). Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. *Management Science*, **41**, 216-232.

Glasserman, P. & S. Tayur (1996). A simple approximation for a multistage capacitated production-inventory systems. *Naval Research Logistics*, **43**, 41-58.

Graves, S. (1985). Multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management Science*, **31**, 1247-1256.

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

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

Graves, S. C. & S. P. Willems (2002). Supply chain design: safety stock placement and supply chain configuration. Chapter 3 in *Handbooks in Operations Research and Management Science*, Vol. 11: Supply Chain Management: Design, Coordination and Operation. A. G. de Kok and S. C. Graves, eds. North-Holland Publishing Company, Amsterdam, The Netherlands.

Heyman, D. & M. Sobel (1984). Stochastic models in operations research. Vol. 1, McGrow-Hill, New York.

Jackson, P. L. (1988). Stock allocation in a two-echelon distribution system or "what to do until your ship comes in". *Management Science*, **34**, 880-895.

Janakiraman, G. & J. A. Muckstadt (2004). A decomposition approach for a class of capacitated serial systems. Working Paper, Stern School of Business. New York University.

Kapuscinski, R. & S. Tayur (1999). Optimal policies and simulation-based optimization for capacitated production inventory systems. In S. Tayur, R. Ganeshan, M. J. Magazine (eds.). *Quantitative Models for Supply Chain Management*, Kluwer Academic Publishers.

Kulkarni, V. G. (1995). Modeling and analysis of stochastic systems. Chapman & Hall, New York.

Lee, H. L. & C. Billington (1993). Material management in decentralized supply chains. *Operations*

Research, **41**, 835-847.

Lee, H. L. & K. Moinzadeh (1987a). Two-parameter approximations for multi-echelon repairable inventory models with batch ordering policy. *IIE Transaction*, **19**, 140-149

Lee, H. L. & K. Moinzadeh (1987b). Operating characteristics of a two-echelon inventory system for repairable and consumable items under batch ordering and shipment policy. *Naval Research Logistics Quarterly*, **34**, 365-380.

Lee, Y. & P. Zipkin (1992). Tandem queues with planned inventories. *Operations Research*, **40**, 936-947.

Lin, G., M. Ettl, S. Buckley, S. Bagchi, D. D. Yao, B. L. Naccarato, R. Allan, K. Kim & L. Koenig (2000). Extended-Enterprise Supply-Chain Management at IBM Personal Systems Group and Other Divisions. *Interfaces*, **30**, 7-25.

Liu, L., X. Liu, D. D. Yao (2004) Analysis and optimization of a multistage inventory-queue system. *Management Science*, **50**, 365-380.

Lu, Y. & J. S. Song (2004). Order-based cost optimization in Assemble-to-order systems. Forthcoming in *Operations Research*.

Muharremoglu, A. & J. N. Tsitsiklis (2002). Echelon base-stock policies in uncapacitated serial inventory systems. Working Paper. Operations Research Center, Massachusetts Institute of Technology.

Muckstadt, J. A. (1973). A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science*, **20**, 472-481.

Muckstadt, J. A. (2004). Analysis and algorithms for service parts supply chains. Springer, New York.

Palm, C. (1938). Analysis of the Erlang traffic formula of busy-signal arrangements. *Ericsson Technics*, **5**, 39-58.

Parker, R. P. & R. Kapuscinski (2004). Optimal policies for a capacitated two-echelon inventory systems. *Operations Research*, **52**, 739-755.

Porteus, E. L. (2002). Foundations of Stochastic Inventory Theory. Stanford University Press.

Rosling, K. (1989). Optimal inventory policies for assembly systems under random demand. *Operations Research*, **37**, 565-579.

Schmidt, C. P. & S. Nahmias (1985). Optimal policy for a two-stage assembly system under random demand. *Operations Research*, **33**, 1130-1145.

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

Sherbrooke, C. (1986). VARI-METRIC: Improved approximations for multi-indenture, multi-echelon availability models. *Operations Research*, **34**, 311-319.

Sigman, K. (2001). Lecture notes on $GI/M/1$ queue. Columbia University. New York, NY.

Simchi-Levi, D. & Y. Zhao (2005). Safety-stock positioning in supply chains with stochastic lead-times. To appear in *Manufacturing and Service Operations Management*.

Song, J. S. (1998). On the order fill rate in a multi-item, base-stock inventory system. *Operations Research*, **46**, 831-845.

Song, J. S. (2000). A note on Assemble-to-Order systems with batch ordering. *Management Sciences*, **46**, 739-943.

Song, J. S. (2002). Order-Based backorders and their implications in multi-item inventory systems. *Management Science*, **48**, 499-516.

Song, J. S. & D. D. Yao (2002). Performance analysis and optimization in Assemble-to-order systems with random lead-times. *Operations Research*, **50**, 889-903.

Song, J. S. & P. Zipkin (2002). Supply chain operations: Assemble-to-order systems. Chapter 11 in *Handbooks in Operations Research and Management Science*, Vol. 11: Supply Chain Management: Design, Coordination and Operation. A. G. de Kok and S. C. Graves, eds. North-Holland Publishing Company, Amsterdam, The Netherlands.

Svoronos, A. & P. Zipkin (1988). Estimating the performance of multi-level inventory systems. *Operations Research*, **36**, 57-72.

Svoronos, A. & P. Zipkin (1991). Evaluation of one-for-one replenishment policies for multiechelon inventory systems. *Management Science*, **37**, 68-83.

Zhao, Y. & D. Simchi-Levi (2005). Performance analysis and evaluation of assemble-to-order systems with stochastic sequential lead-times. To appear in *Operations Research*.

Zipkin, P. (1991). Evaluation of base-stock policies in multiechelon inventory systems with compound-Poisson demands. *Naval Research Logistics*, **38**, 397-412.

Zipkin, P (2000). Foundations of Inventory Management. McGraw Hill, Boston.

Appendix I

Table 10: Solutions for example 2

node	$n = 4$	$n = 9$	$n = 16$
Platform Group	20	5	5
Fender Group	23	9	9
Roll Over Group	5	0	0
Chassis/Platform	80	3	1
Frame Assembly	69	69	60
Case	0	0	0
Case & Frame	141	128	125
Brake Group	0	0	0
Drive Group	18	13	9
Plant Carrier	29	17	9
Final Drive & Brake	38	36	36
Engine	0	0	0
Fans	38	30	21
Dressed-out Engine	94	2	2
Boggie Assembly	17	13	9
Pin Assembly	259	206	186
Suspension Group	80	73	64
Transmission	0	0	0
Common Subassembly	90	1	5
Main assembly	88	55	146
Track Roller Frame	58	50	7
Final Assembly	159	139	130