Numerical investigation of tradeoffs in production-inventory control policies with advance demand information

George Liberopoulos and Stelios Koukoumialos

University of Thessaly, Department of Mechanical and Industrial Engineering, Volos, Greece, Tel: +30 421 74056, E-mails: glib@mie.uth.gr and skoukoum@mie.uth.gr

Abstract

This paper investigates tradeoffs between optimal base-stock levels, numbers of kanbans, and lead-times in base-stock and extended kanban policies with advance demand information used for production-inventory control in multi-stage supply chains. Simulation-based computational experience regarding such tradeoffs is reported for single-stage and two-stage production-inventory supply chains.

Keywords: *Production-inventory control, advance demand information, lead-time, base-stock, kanban*

1. Introduction

In a multi-stage production-inventory supply chain, end-item demand information may be available in advance in the form of actual orders, commitments, forecasts, etc., to all stages of the supply chain. Recent developments in information technology and the emphasis on supply chain integration have significantly reduced the cost of diffusing advance demand information to all stages of the supply chain and have created an opportunity for the development of production-inventory control policies that use such information. The implementation of such policies may result in significant cost savings throughout the entire supply chain through inventory reductions as well as improvements in customer service.

In this paper we consider policies that use advance demand information for the production-inventory control of a multi-stage serial supply chain that produces a single type of job in a make-to-stock mode. Every stage in the supply chain consists of a facility where jobs are processed, and an output store where finished jobs are stored. Jobs in the facility are called *work-in-process* (WIP), and jobs in the output store are called *finished goods* (FG). FG of the last stage are called *end-items*. There is an infinite supply of raw jobs feeding the first stage. Customer demands arrive randomly for one end-item at a time, with a constant demand lead-time in advance of their due-dates. Once a customer demand arrives, it can not be cancelled, i.e. the advance demand information is assumed to be perfect. Demands that are not satisfied from FG inventory on their due-dates are backordered and are called *backordered demands* (BD). The arrival of a customer demand for an end-item triggers a replenishment order for FG inventory at every stage. FG inventory levels at all stages are followed continuously, and replenishments of FG inventory may be ordered at any time. There is no setup cost or setup time for placing an order and no limit on the number of orders that can be placed per unit time. Under the above assumptions, there is no incentive to replenish FG inventory by anything other than a continuous review, one-for-one replenishment policy.

When there is no advance demand information, demand due-dates coincide with demand arrival times. In this case, the replenishment orders at every stage, triggered by the arrival of customer demands, are issued at or after the demand due-dates. The simplest production-inventory control policy, when there is no advance demand information, is a *base-stock policy*. Base-stock policies were originally designed for non-capacitated inventory systems. A policy that has attracted considerable attention and has particular appeal in a JIT capacitated production environment is a *kanban policy*, which in the case of a single-stage system is equivalent to a make-to-stock *CONWIP policy* [8]. Base-stock and kanban policies may be combined to form more sophisticated hybrid *base-stock/kanban policies* [8].

When there is advance demand information, the replenishment orders for FG inventory at every stage, triggered by the arrival of a customer demand, may be issued before the due-date of the demand. Base-stock and base-stock/kanban policies can be easily modified to take advantage of advance demand information by offsetting demand due-dates by stage *lead-times* to determine the issue times of replenishment orders at every stage. A kanban policy can not exploit advance demand information because in a kanban policy an order is issued after a job in FG inventory is consumed and therefore at or after the due-date of the demand that triggered it.

The aim of this paper is to investigate the tradeoffs between optimal base-stock levels, numbers of kanbans, and lead-times in multi-stage base-stock and base-stock/kanban policies, and in particular extended kanban policies, with advance demand information. In the sections that follow we report on simulation-based computational experience regarding these tradeoffs in single-stage and two-stage supply chains.

2. Single-stage base-stock policy with advance demand information

In this section we consider a single-stage base-stock policy with advance demand information similar to that considered in [4]. Customer demands arrive for one end-item at a time according to a Poisson process with rate λ, with a constant demand lead-time, T, in advance of their due-dates. The arrival of every customer demand triggers the consumption of an end-item from FG inventory and the issuing of a replenishment order to the facility of the single stage. The consumption of an end-item from FG inventory is triggered T time units after the arrival time. If no end-items are available at that time, the demand is backordered. The system starts with a base-stock of S end-items in FG inventory. The time of issuing the replenishment order is determined by offsetting the demand due-date by the stage lead-time, L. This means that the order is issued with no delay, if $L \geq T$, or with a delay of $T - L$ with respect to the demand arrival time, if $L < T$. In short, the delay in issuing an order is equal to $\max[0, T - L]$. When the order is issued, a new job is immediately released into the facility. If there is no advance demand information, i.e. if $T = 0$, both the consumption of an end-item from FG inventory and the replenishment order are triggered at the demand arrival time and the resulting policy is a classical base-stock policy.

A queuing network model of a base-stock policy with advance demand information is shown in Figure 1. The symbolism used in Figure 1 (and in other similar figures that follow) is the same as that used in [4] and [8] and has the following interpretation. The oval represents the facility, and the circles represent time delays. The queues followed by vertical bars represent *synchronization stations* linking the queues. A synchronization station is a server with instant service time that "fires" as soon as there is at least one "customer" in each of the queues that it synchronizes. Queues are labeled according to their content, and their initial value is indicated inside parentheses. Queue OH contains orders on hold. Notice that queue OH is always equal to zero, since there is an infinite number of raw jobs.

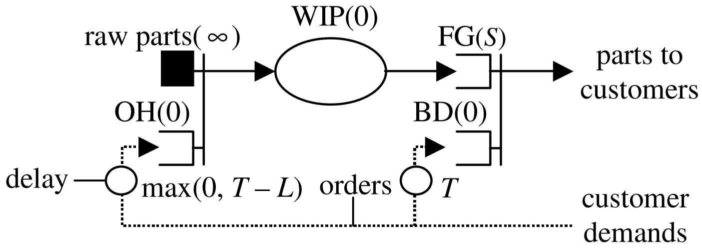

Figure 1. Single-stage base-stock policy with advance demand information

We consider a typical optimization problem where the objective is to find the values of S and L that minimize the long-run expected average cost of holding and backordering inventory,

$$C(S,L) = hE[WIP + FG(S,L)] + bE[BD(S,L)], \tag{1}$$

where h is the unit cost of holding WIP + FG inventory per unit time and b is the unit cost of backordering FG inventory per unit time. Notice that control parameters S and L affect the expected average FG and BD only and not the expected average WIP.

If $T = 0$, L is irrelevant, and the optimal base-stock level, S^*, is the smallest integer that satisfies [10]

$$P(WIP \le S^*) \ge b/(b+h). \tag{2}$$

If the facility consists of a Jackson network of servers, S^* satisfies a non-closed-form expression that can be solved numerically [10].

If $T > 0$ and the facility consists of a single-server station with exponential service rate μ, the optimal lead-time L^* does not depend on T. Moreover, $S^* = \lfloor \hat{S} \rfloor$, where $\hat{S}$ is a function of T that decreases linearly with T and reaches zero at $T = L^*$ (sec. 4.5.2) [1], [6]. Thus, for $T < L^*$, $\hat{S} > 0$ and orders are issued upon the arrival of demands with no delay, whereas for $T > L^*$, $\hat{S} = 0$ and orders are issued with a delay of $T - L^*$ with respect to the arrival times of demands. Since L^* is the smallest value of T for which $\hat{S}$ = 0, and $S^* = \lfloor \hat{S} \rfloor$, the smallest value of T for which $S^* = 0$, is just below L^*.

If the facility consists of a Jackson network of servers, there are no analytical results for the optimal parameter values. To shed some light into this case, we numerically investigated a particular instance of the system in which the facility consists of a Jackson network of $M = 4$ identical single-server stations in series, each server having an exponential service rate μ. For this instance, we considered two sets of parameter values shown in Table 1. We used simulation to evaluate the long-run expected average cost of the system for the two sets of parameter values. In each case we optimized the control parameters S and L for different values of T, using exhaustive search. The optimization yielded the following results.

Table 1. Parameter values for cases 1 and 2 of a single-stage base-stock policy

Case	$1/\lambda$	$1/\mu$	$\rho = \lambda/\mu$	H	b
1	1.25	1.0	0.8	5	1
2	1.1	1.0	0.90909…	1	9

For $T = 0$, L is irrelevant, and S^* can be determined analytically [10]. The results are $S^* = 8$ and $S^* = 68$, for cases 1 and 2, respectively. As T increases away from zero, S^* decreases (apparently linearly) with T and reaches zero just below $T = L^*$, as in the case of the single-server station. Plots of S^* versus T are shown in Figure 2 for both cases.

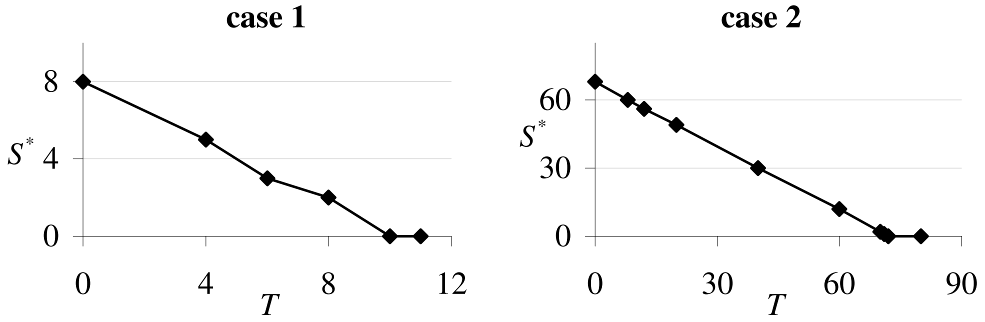

Figure 2. S^* **versus** T **for a single-stage base-stock policy**

From Figure 2 it can be seen that the smallest values of T for which $S^* = 0$ are approximately equal to 10 and 72, for cases 1 and 2, respectively. We say "approximately" because we only examined integer values of T, whereas T really is a continuous parameter. The optimal lead-times L^* are slightly larger than 10 and 72, for cases 1 and 2, respectively. The minimum long-run expected average cost $C(S^*,L^*)$ decreases *very little* with T and attains its minimum value at $T = L^*$.

3. Single-stage extended kanban policy with advance demand information

A single-stage extended kanban policy with advance demand information behaves similarly to a single-stage base-stock policy with advance demand information. The only difference is that in a single-stage extended kanban policy, when an order is issued, it is not immediately authorized to go through, unless the inventory in the system, i.e. WIP + FG, is below a given *inventory-limit* of K jobs. If the inventory in the system is at or above K, the order is put on hold until the inventory drops below K (the inventory drops as FG end-items are consumed by customers). Once the order is authorized to go through, a new job is immediately released into the facility. This policy can be implemented by requiring that every job entering the facility be granted a production authorization card called *kanban*, where the total number of kanbans is equal to the constant inventory level. Once a job leaves the FG output store, the kanban that was granted (and attached) to it is detached and is used to authorize the release of a new job into the facility. The system starts with a base-stock of S end-items in FG inventory and $K - S$ free kanbans that are available to authorize an equal number of orders. This means that K and S must satisfy

$$S \leq K . \tag{3}$$

The number of free kanbans represents the number of jobs that can be released into the facility before the inventory in the system reaches the inventory-limit K. A queueing network model of a base-stock policy with advance demand information is shown in Figure 3, where queue FK contains *free kanbans*.

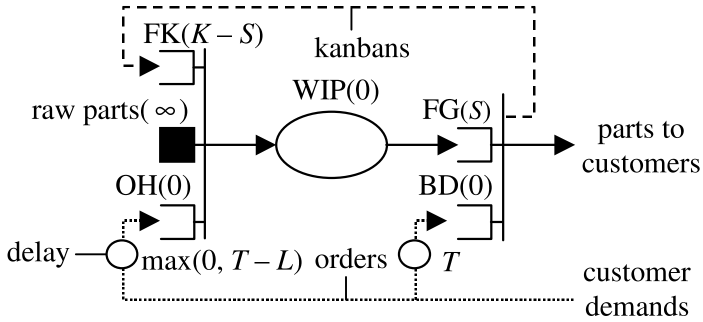

Figure 3. Single-stage extended kanban policy with advance demand information

From Figure 3 it can be seen that kanbans trace a loop within a closed network linking FK, WIP, and FG. The constant population of this closed network is K, i.e. at all times, FK + WIP + FG = K. The throughput of this closed network, denoted TH_K, depends on K and determines the processing capacity of the system, i.e. the maximum demand rate λ that the system can meet in the long run. Under some fairly general conditions, for every feasible demand rate λ, such $\lambda < \mathrm{TH}_\infty$, there is a finite minimum value of K, $K_{\min}$, such that for any $K \geq K_{\min}$, $\mathrm{TH}_K > \lambda$, which means that the system has enough capacity to meet demand in the long run. Notice that a single-stage extended kanban policy with $K = \infty$ and $S < \infty$ is equivalent to a single-stage base-stock policy with the same S. A single-stage extended kanban policy with $K = S < \infty$ is equivalent to a single-stage kanban policy or make-to-stock CONWIP policy with K kanbans.

We consider an optimization problem similar to that in Section 2, where the objective is to find the values of K, S, and L that minimize the long-run expected average cost of holding and backordering inventory,

$$C(K,S,L) = hE[WIP_K + FG_K(S,L)] + bE[BD_K(S,L)], \quad (4)$$

where h and b are defined as in Section 2. Notice that control parameters S and L affect the expected average FG and BD only and not the expected average WIP or OH, whereas parameter K affects the expected average FG, BD as well as WIP and OH.

If T = 0, the lead-time parameter L is irrelevant, and the optimal base-stock level for any given value of K, such that $K \geq K_{\min}$, S_K^*, is given by $S_K^* = \min[K, S_K^{**}]$, where S_K^{**} is the smallest integer that satisfies [9]

$$P(OH_K + WIP_K \leq S_K^{**}) \geq b/(b+h). \quad (5)$$

If the facility consists of a Jackson network of servers, there is no analytical expression (not even a non-closed-form one) to determine the steady-state distribution of OH_K and WIP_K and therefore S^*, and only approximate methods exist [7]. To shed some light into this case, we numerically investigated the same instance of the system as that in Section 2 for the same two sets of parameter values shown in Table 1. For this instance, $\mathrm{TH}_K = \mu/[1 + (M - 1)/K]$ [7], which implies that $K_{\min}$ is the smallest integer that satisfies $K_{\min} > (M - 1)\rho/(1 - \rho)$ and is equal to 13 and 30, for cases 1 and 2, respectively. We used simulation to evaluate the long-run expected average cost of the system for the two sets of parameter values, and in each case we found the optimal unconstrained and constrained base-stock levels for different values of K, S_K^{**} and S_K^{*}, using exhaustive search. Plots of S_K^{**} and S_K^{*} versus K are shown in Figure 4 for both cases. Based on the results shown in Figure 4, we make the following observations.

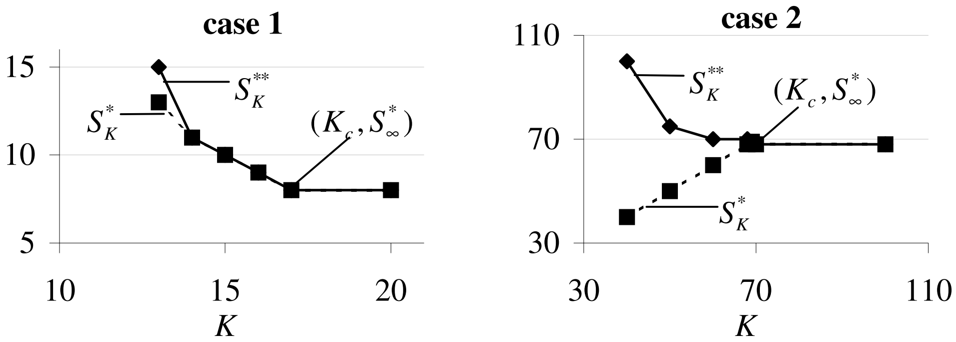

Figure 4. S_K^{**} and S_K^{*} versus K for a single-stage extended kanban policy

The unconstrained optimal base-stock level S_K^{**} is non-increasing in K, i.e. $S_{K+1}^{**} \le S_K^{**}$, for $K \ge K_{\min}$. Moreover, there exists a finite critical value of K, K_c, such that $S_K^{**} = S_\infty^*$, for $K \ge K_c$, where S_∞^* is the optimal base-stock level for the same system operating under a pure base-stock policy. This means that there is a trade-off between K and S_K^{**} and that this trade-off holds for up to a finite critical value of K, K_c. This critical value is equal to 17 and 70, for cases 1 and 2, respectively. This result is shown analytically in [9] for a slightly different but equivalent system. The intuition behind it is the following. Increasing K increases the processing capacity of the facility and therefore the average FG inventory. Increasing K also increases the congestion in the facility. At $K = K_c$, the facility reaches a critical congestion level. That is, for values of K below K_c, the system is under-congested in the sense that increasing K increases the average FG inventory enough to warrant a decrease in S_K^{**}. For values of K above K_c, however, the system is over-congested in the sense that increasing K does not increase the average FG inventory enough to warrant a decrease in S_K^{**}.

The most interesting result of the optimization is that the overall optimal number of kanbans K^* is equal to 17 and 70, for cases 1 and 2, respectively, and the overall optimal base-stock level S^* is equal to 8 and 68, for cases 1 and 2, respectively. In other words, $K^* = \max[K_c, S_\infty^*]$ and $S^* = S_{K^*}^* = S_\infty^*$. This is not an obvious result. It means that S^* is equal to the smallest possible value of S_K^{**}, i.e. S_∞^*, and that K^* is the smallest value of K for which $S_K^{**} = S_K^* = S_\infty^*$. The intuition behind this result is that it is optimal to set K to a value that is just big enough so that the corresponding unconstrained and constrained optimal base-stock levels are equal to S_∞^*. Computational experience reported in [3], [5], and [11] (sec. 8.8.2) also validates this result. The difficulty in proving it stems from the fact that no analytical expression for the steady-state distribution of OH_K and WIP_K exists, except for a trivial system where the facility consists of a single-server station with exponential service rate μ, in which case $K_{\min} = K_c = 1$. Nevertheless, a strong indication of the validity of this result is given by Proposition 6 in [9].

If $T > 0$, and the facility consists of a Jackson network of servers, there are no analytical results available for the optimal parameter values. To shed some light into this case, we numerically investigated the same instance of the system as that in the case where $T = 0$. We used simulation to evaluate the long-run expected average cost of the system for the same two sets of parameter values, and in each case we optimized the control parameters, K, S and L, for different values of T, using exhaustive search. The optimization yielded the following results.

The optimal number of kanbans, K^*, is equal to $\max[K_c, S_\infty^*]$ for all values of T, i.e. it is independent of T. Moreover, L^* and S^* have the same values as those in the single-stage base-stock policy with advance demand information discussed in Section 2. The intuition behind this result is the following. The optimal value of K determines the optimal processing capacity and congestion level in the facility. T does not affect the facility. Therefore, the value of K that is optimal for $T = 0$ is also optimal for $T > 0$. T only affects FG and BD and as such it is a substitute for S^*, which also affects only FG and BD. Since for $T = 0$, $S^* = S_\infty^*$, for $T > 0$, the tradeoff between T and S^* is exactly the same as the tradeoff between T and S_∞^*. Therefore, as T increases away from zero, S^* decreases and reaches zero at T just below L^*, exactly as in a pure base-stock policy.

4. Two-stage base-stock policy with advance demand information

In this section we consider a two-stage base-stock policy with advance demand information similar to that considered in [4]. Customer demands arrive for one end item at a time according to a Poisson process with rate λ, with a constant demand lead-time, T, in advance of their due-dates. The arrival of every customer demand triggers the consumption of an end-item from FG inventory and the issuing of replenishment orders to the facilities of each of the two stages in the supply chain. The consumption of an end-item from FG inventory is triggered T time units after the arrival time. If no end-items are available at that time, the demand is backordered. The supply chain starts with a base-stock of S_1 and S_2 jobs in FG inventory at stages 1 and 2, respectively. The time of issuing the replenishment order at stage 2 is determined by offsetting the demand due-date by the stage lead-time, L_2. The time of issuing the replenishment order at stage 1 is determined by offsetting the demand due-date by the sum of the lead-times of stages 1 and 2, $L_1 + L_2$. As in the single-stage case, this means that the delay in issuing an order at stage 2 is equal to $\max[0, T - L_2]$, whereas the delay in issuing an order at stage 1 is equal to $\max[0, T - (L_1 + L_2)]$. When the order is issued at stage 1, a new job is immediately released into the facility of stage 1. When the order is issued at stage 2, a new job is also immediately released into the facility of stage 2, provided that such a job is available in the FG output store of stage 1. Otherwise, the job remains on hold until a job becomes available in the FG output store of stage 1. If there is no advance demand information, i.e. if $T = 0$, both the consumption of an end-item from FG inventory and the replenishment orders are triggered at the demand arrival time, and the resulting policy is a classical base-stock policy. A queuing network model of a base-stock policy with advance demand information is shown in Figure 5.

We consider an optimization problem similar to that in Section 2, where the objective is to find the values of S_1, S_2, L_1, and L_2 that minimize the long-run expected average cost of holding and backordering inventory,

$$
\begin{aligned}
C(S_1,S_2,L_1,L_2) = {} & h_1 E[WIP_1 + FG_1(S_1,L_1,L_2)] \\
& + h_2 E[WIP_2(S_1,L_1,L_2) + FG_2(S_1,S_2,L_1,L_2)] + bE[BD(S_1,S_2,L_1,L_2)]
\end{aligned} \tag{6}
$$

where h_n is the unit cost of holding WIP + FG inventory per unit time at stage n and b is the unit cost of backordering end-item inventory per unit time.

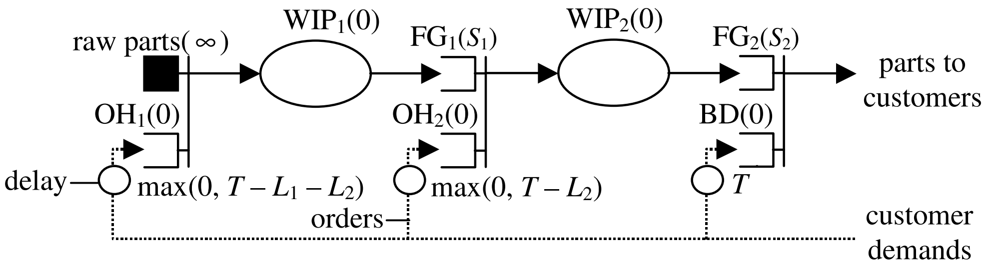

Figure 5. Two-stage base-stock policy with advance demand information

If $T = 0$, the lead-time parameters L_1 and L_2 are irrelevant. Unfortunately, there are no analytical results available for the optimal base-stock levels, S_1^* and S_2^*, even in the case where each facility consists of a Jackson network of servers. Some approximation methods have been developed in [1] (sec. 10.7), [3], and [11] (sec. 8.3.4.3). The only analytically tractable case is the case where $S_1 = 0$. In this case, the two-stage policy is equivalent to a single-stage policy where the facilities of stages 1 and 2 are merged into a single facility. This is useful to know because in case $h_1 \geq h_2$ it is not too difficult to see that $S_1^* = 0$. This means that the only interesting case is the case where $h_1 < h_2$.

If $T > 0$, there are no analytical results available for the optimal parameter values. To shed some light into this case, we numerically investigated a particular instance of the system in which each facility consists of a Jackson network of $M = 2$ identical single-server stations in series, each server having an exponential service rate μ. For this instance, we considered the set of parameter values shown in Table 2.

Table 2. Parameter values for the case of a two-stage base-stock policy

Case	$1/\lambda$	$1/\mu$	$P = \lambda/\mu$	h_1	h_2	b
1	1.1	1.0	0.90909…	1	3	9

For this set of parameter values, we used simulation to evaluate the long-run expected average cost of the system, and we optimized the control parameters S_1, S_2, L_1, and L_2 for different values of T, using exhaustive search. The optimization yielded the following results.

For $T = 0$, L_1 and L_2 are irrelevant and $S_1^* = 26$ and $S_2^* = 31$. As T increases away from zero, S_1^* remains constant, while S_2^* decreases (apparently linearly) with T and reaches zero just below $T = L_2^*$. As T increases away from L_2^*, S_2^* remains zero, while S_1^* decreases (apparently linearly) with T and reaches zero just below $T = L_1^* + L_2^*$. Plots of S_1^* and S_2^* versus T are shown in Figure 6. The intuition behind this result is the following.

When $T = 0$, $S_1^* > 0$, $S_2^* > 0$, and orders are issued at both stages with no delay. As T increases away from zero, it is optimal to reduce only S_2^* and not S_1^*, because $h_2 > h_1$. When T is just below L_2^*, S_2^* becomes zero. As T increases away from L_2^*, S_2^*

remains zero, and orders are issued at stage 2 with a delay of $T - L_2^*$. At the same time, S_1^* starts decreasing with *T*, and orders are still issued at stage 1 with no delay. When *T* is just below $L_1^* + L_2^*$, S_1^* becomes zero too. As *T* increases away from $L_1^* + L_2^*$, both S_1^* and S_2^* remain at zero, while orders are issued at stages 2 and 1 with delays of $T - L_2^*$ and $T - (L_1^* + L_2^*)$, respectively, i.e. in each stage, orders are issued with a delay only when *T* is large enough so that the optimal *echelon* base-stock level of the stage is zero.

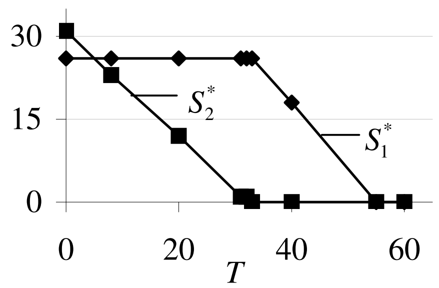

Figure 6. S_1^* **and** S_2^* **versus** ***T*** **for a two-stage base-stock policy.**

From Figure 6 it can be seen that the smallest values of *T* for which $S_2^* = 0$ and $S_1^* = 0$ are approximately equal to 33 and 55, respectively. The optimal lead-times L_2^* and L_1^* are a bit larger than 33 and 22. As in the case of the single-server station, the optimal lead-times are independent of *T*. The minimum long-run expected average cost decreases *very little* with *T* and attains its minimum value at $T = L_1^* + L_2^*$.

5. Two-stage extended kanban policy with advance demand information

A two-stage extended kanban policy with advance demand information is a straightforward extension of a single-stage extended kanban policy with advance demand information presented in Section 3 to two stages. Due to space considerations, we do not include a detailed description of it here. A queuing network model of an extended kanban policy with advance demand information is shown in Figure 7.

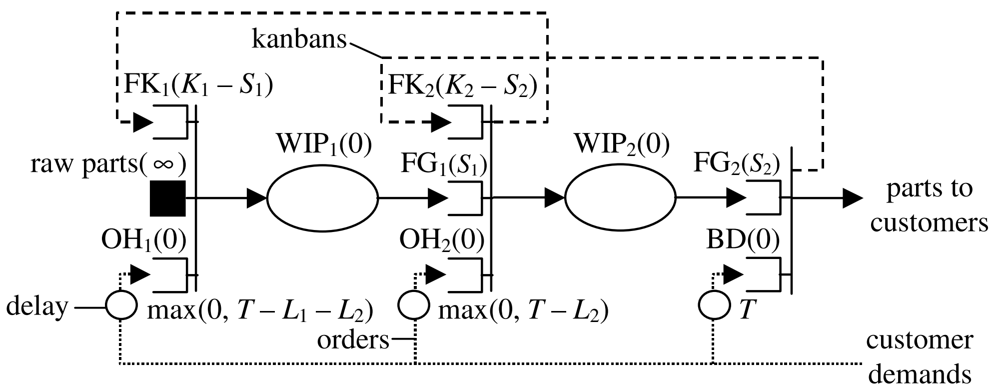

Figure 7. Two-stage extended kanban policy with advance demand information

If $T > 0$, there are no analytical results available for the optimal parameter values. To shed some light into this case, we numerically investigated the same instance of the system as that in Section 4, for the same set of parameter values shown in Table 2. For this set of parameter values, we used simulation to evaluate the long-run expected average cost of the system, and we set to optimize the control parameters K_1, K_2, S_1, S_2, L_1, and L_2 for different values of T, using exhaustive search. The preliminary results of the optimization indicate that the properties of the optimal parameter values are similar to those of the optimal parameter values in the single-stage extended kanban policy.

Namely, for $T = 0$, L_1 and L_2 are irrelevant, and S_1^* and S_2^* are equal to the optimal base-stock levels for the two-stage pure base-stock policy, i.e. $S_1^* = 26$ and $S_2^* = 31$. Moreover, the optimal numbers of kanbans K_1^* and K_2^* are the smallest values of K_1 and K_2 for which the optimal base-stock levels are equal to the optimal base-stock levels in the two-stage pure base-stock policy. These values are $K_1^* = 42$ and $K_2^* = 31$.

For $T > 0$, K_1^* and K_2^* remain constant for all values of T, whereas L_1^*, L_2^*, S_1^*, and S_2^* have the exact same values as in the two-stage base-stock policy with advance demand information discussed in Section 4. The intuition behind these preliminary results is the same as that behind the results for the single-stage extended kanban policy.

References

[1] Buzacott, J.A. and Shanthikumar J.G. (1993) *Stochastic Models of Manufacturing Systems*, Prentice-Hall, Englewood Cliffs, NJ.

[2] Dallery, Y. and Liberopoulos, G. (2000) Extended kanban control system: Combining kanban and base stock, IIE Transactions on Design and Manufacturing, 32 (4), 369-386.

[3] Duri, C., Frein, Y., and Di Mascolo, M. (2000) Comparison among three pull control policies: kanban, base stock and generalized kanban, Annals of Operations Research, 93, 41-69.

[4] Karaesmen, F., Buzacott, J.A., and Dallery, Y. (1999) Integrating advance order information in make-to-stock production systems, working paper, Laboratoire Productique-Logistique, Ecole Centrale de Paris.

[5] Karaesmen, F. and Dallery, Y. (2000) A performance comparison of pull control mechanisms for multi-stage manufacturing systems, International Journal of production Economics, 68, 59-71.

[6] Karaesmen, F., Liberopoulos, G., and Dallery, Y. (2000) The value of advance information on demands for make-to-stock manufacturing systems, working paper, Laboratoire Productique-Logistique, Ecole Centrale de Paris.

[7] Frein, Y., Di Mascolo, M. and Dallery, Y. (1995) On the design of generalized kanban control systems, International Journal of Operations and Production Management, 15 (9), 158-184.

[8] Liberopoulos, G. and Dallery, Y. (2000) A unified framework for pull control mechanisms in multi-stage manufacturing systems, Annals of Operations Research, 93, 325-355.

[9] Liberopoulos, G. and Dallery, Y. (2000) Base-stock versus WIP-limit in single-stage make-to-stock production-inventory systems, working paper, Department of Mechanical and Industrial Engineering, University of Thessaly, Greece.

[10] Rubio R. and Wein, L.W. (1996) Setting base stock levels using product-form queueing networks, Management Science, 42 (2), 259-268.

[11] Zipkin, P. (2000) *Foundations of Inventory Management*, McGraw Hill: Management & Organization Series, Boston, MA.