JOURNAL OF TEXTILE AND APPAREL, TECHNOLOGY AND MANAGEMENT

Volume 2, Issue 2, Spring 2002

Control of Yarn Inventory for a Cotton Spinning Plant: Part 2: Correlated Demand and Seasonality

Russell E. King[1]
Leigh Ann C. Brain[1]
Kristin A. Thoney[2]

[1]Department of Industrial Engineering
[2]Department of Textile and Apparel, Technology and Management
North Carolina State University

This research was supported, in part, by the National Textile Center through the U.S. Department of Commerce.

ABSTRACT

In the first paper, a novel application of process control was developed for controlling inventory in a yarn spinning plant. This technique, proportional/derivative control, was compared to a traditional min/max system via computer simulation for situations in which demand is uncorrelated. In this second paper, the performance of both techniques is compared in situations in which the demand from period to period exhibits correlation. In addition, demand seasonality is also considered. The proposed methodology is shown to provide superior control allowing lower levels of inventory with fewer frame changeovers to achieve the same levels of shipping performance. It also is able to quickly adjust inventories to track changes in demand.

KEYWORDS: spinning, proportional derivative control, min max control, inventory control, process control, cotton

1. Introduction

In Part 1, the simulation model assumed basic goods with an independently and identically distributed (i.i.d.) stochastic demand for contract orders generated from a normal distribution with a constant mean and variance for each yarn type *i*. Experimentation under this model showed the discrete PD algorithm to be superior to the min/max algorithms. However, typically real-world spinning plants do not experience uncorrelated demand. With this in mind, the question arises of how the performance of the algorithms is affected by correlated demand. In addition, to validate fully the appropriateness of the PD algorithm, it is also tested using seasonal goods which create non-stationary conditions.

2. Correlated Demand

2.1 Generation of Correlated Demand in the Model

Correlated demand was created using a time series moving average lag 1 model, MA(1). This implies that the current demand is influenced by the previous demand.

$$W_{t+1,i} = u_i + \theta e_{t,i} + e_{t+1,i}$$

where

$W_{t+1,i}$ = the demand for yarn type *i* for week *t+1,*

u_i = the average demand for yarn type *i,*

θ = the demand correlation, and

$e_{t,i}$ = the deviation of the demand in week *t* from the average demand for yarn *i*, $\{e_{t,i} : t = 0,1,2,... \}$ *i.i.d.* $N(0, \sigma_e^2)$.

In the simulation, the correlation factor was set at 0.75. The weekly error for yarn type *i* was generated by sampling from a normal distribution with a zero mean and standard deviation equal to the demand standard deviation in the uncorrelated demand scenario multiplied by the percentage of demand for yarn type *i* (p_i). The average demand for yarn type *i* was generated by multiplying the average demand in the uncorrelated demand scenario, λ, by the percentage of demand for yarn type *i* (p_i). The output of the MA(1) formula yields the number of pounds demanded in week *t* for yarn type *i*. This value is then divided by the average order size and rounded to the nearest integer to determine the number of orders.

$$o_{t,i} = \left[\frac{W_{t,i}}{\overline{O}} \right]$$

where

$o_{t,i}$ = the number of orders for yarn type *i* for week *t,*

$W_{t,i}$ = the demand in lbs for yarn type *i* for week *t,* and

$\overline{O}$ = the average customer order size.

2.2 PD Algorithm and Correlated Demand

The optimal parameters for c_1 and c_2 that were determined using the uncorrelated demand may not hold under the correlated demand scenario. Thus, the steps described in Part 1 were repeated.

Using the same experimental design as in the uncorrelated scenarios, sixteen runs were made at each of 36 design points where a design point is a (c_1, c_2) combination. The resulting data points show a general minimum value around c_1 = 3 and c_2 = 0.1. These values were then used as the starting point for the Nelder-Mead optimization procedure. The stopping criteria was achieved with simplex 19 when the differences between the largest and smallest c_1 and c_2 values in the simplex was less than 0.001. This experimentation shows the optimal point to indeed be the starting point, c_1 = 3 and c_2 = 0.1. Unlike the uncorrelated demand scenario, c_2 in this case is not near zero. This is because c_2 measures the direction of the error curve, not the error itself. In the uncorrelated demand model, there should not be any trend in the demand and therefore c_2 should not be a significant factor. This is indeed the case as the optimal c_2 value is zero. However, in the correlated demand model, there is a trend in the demand; this week's demand is directly related to last week's demand. Thus, it is expected that the parameter that considers the slope and direction of the error curve would be significant.

2.3 PD vs. Min/Max Algorithm Under Correlated Demand

To determine if the PD algorithm is still superior under correlated demand, it is compared to its best rival from the uncorrelated demand scenario, the min/max algorithm using a 5 day range and no dedicated frames. Each algorithm was run using a full design of four factors with two levels for each factor. Again the target inventory level for the PD algorithm was varied from 1 to 15 days and the range for the min/max algorithm was varied from 1-6 days to 18-22 days. The PD algorithm used the optimal c_1 and c_2 values found in the previous section.

A graph of the percentage of on-time orders for both algorithms is shown in Figure 1. While both algorithms perform similarly at high inventory levels, the PD algorithm outperforms the min/max at lower inventory levels. The inventory standard deviation in Figure 2 is much lower for the PD algorithm. The PD algorithm also achieves roughly 13 fewer frame changeovers per day than the min/max algorithm as seen in Figure 3.

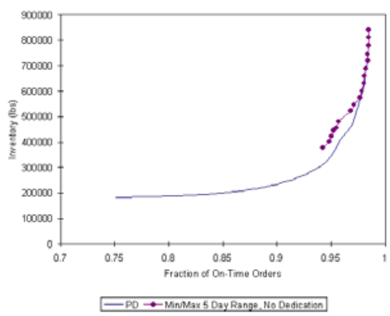

Figure 1: Percentage of On-Time Orders vs. Inventory Under Correlated Demand

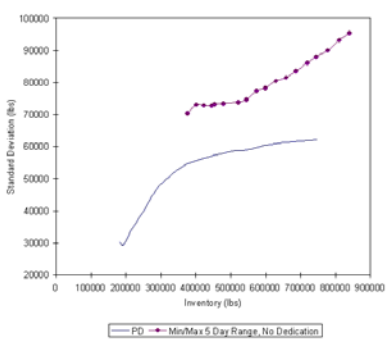

Figure 2: Inventory vs. Inventory Standard Deviation Under Correlated Demand

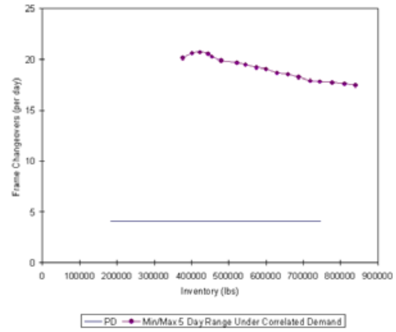

Figure 3: Inventory vs. Frame Changeovers Under Correlated Demand

2.4 Correlated Demand vs. Uncorrelated Demand

In order to assess the impact on key output parameters, the uncorrelated results from Part 1 (**Figures 2-4**) were compared with the correlated demand results (Figures 1-3). While both the PD and min/max algorithms performed better under correlated demand, the PD algorithm saw the most improvement for the percentage of on-time orders and the min/max algorithm experienced a decrease of approximately 3 frame changeovers per day. This matches intuition since the inclusion of the derivative term helps predict how the inventory level is changing. The algorithm is able to increase/decrease production based on the prediction.

2.5 Conclusions

Correlated demand allows both algorithms to react quicker to customer orders achieving the same or better performance outputs than the uncorrelated scenario. In areas where the min/max algorithm outperformed the PD algorithm under uncorrelated demand, the PD algorithm is as good as or superior to the min/max algorithm under correlated demand. Thus in the more real-world scenario of correlated demand, the PD algorithm is still superior to the min/max algorithm with regard to the five key output parameters mentioned.

3. Seasonality in Demand

3.1 Seasonality in the Model

The spinning plant simulation model allows up to 4 seasonal changes. Each season can have a unique duration. The customer demand changes instantaneously at the time of season change. If the model can be shown to handle an instantaneous demand change, it can also handle the simpler and more realistic scenario of a smoother demand change. A moving target level for each yarn type is computed each day starting D_{tar} days before the new season and lasting until the new season begins at which time the new seasonal target level is computed. The moving target level is determined by summing the next D_{tar} days of demand on any given day. This provides a transitional period of demand such that at the start of the new season the actual level of inventory for each yarn corresponds to the new seasonal demand. Figure 4 shows an example for the computation of a moving target with an eight day target level. The target level 5 days before the start of the new season is the average of five times the demand in season 1 multiplied by three times the demand in season 2.

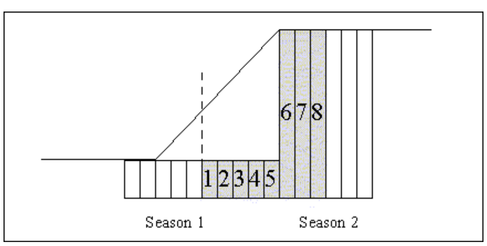

Figure 4: Moving Target Computation for an Eight Day Target Level

3.2 Experimentation

The model was run at three target levels (6 days, 12 days, and 18 days) for five and one quarter years with statistics cleared after the first quarter. A season change was experienced every six months. Five full seasons worth of data was averaged together to create a snapshot of the behavior of the system at the given target NIP level. The purpose of this experimentation is to examine how quickly the model can react to a new higher demand level and its behavior when demand drops back to its original level. **The model was run using an 18 yarn case with extreme (season one) / reverse extreme (season two) demand rates.** Important factors examined are the average overall daily and weekly NIP values as they compare to the target NIP value, and the average weekly order lead-time especially around the time of the seasonal change. Also of interest is the behavior of the yarn types that experience the most extreme seasonal changes. Yarn type #1 has a high production rate. Yarn type #18 has a low production rate. In season one, yarn type #1 has a high demand rate while yarn type #18 has a low demand rate. During season two, the demand rates switch; yarn type #1 is in low demand, while yarn type #18 is in high demand.

3.3 Results

Figures 5, 6, and 7 are graphs of overall inventory levels using 6, 12, and 18 day target NIP values, respectively. Each graph shows the target NIP level, the average daily NIP value, the average weekly NIP value and the average weekly order lead-time. In each case, switching from a lower target level to a higher target level is more difficult to track. The larger the difference between seasonal target levels, the longer it takes the algorithm to reach its new target level. However, the target level does not affect the reactivity when switching from a high demand season to a low demand season simply because production can be halted and orders shipped from inventory. Order lead-times are not significantly affected during the seasonal change.

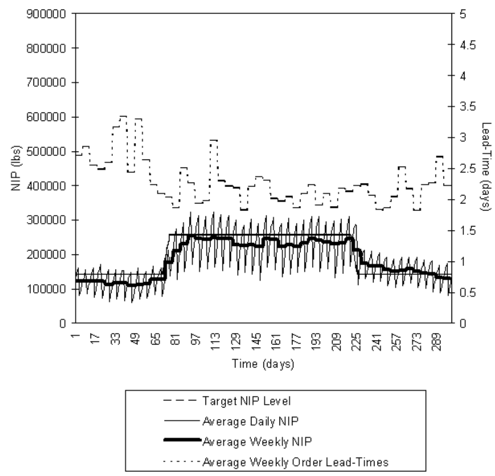

Figure 5: Seasonality Using a 6 Day Target Level

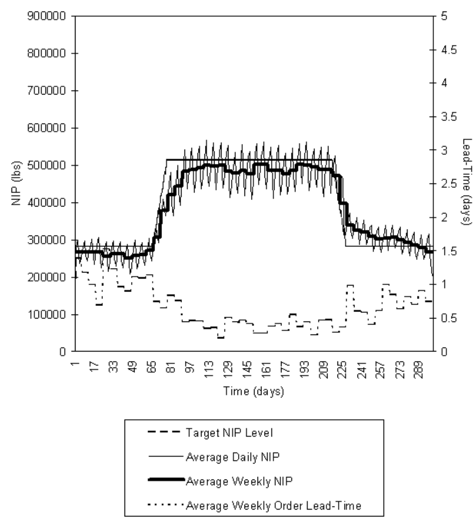

Figure 6: Seasonality Using a 12 Day Target Level

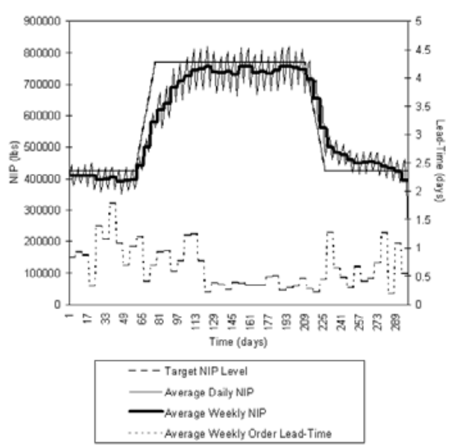

Figure 7: Seasonality Using a 18 Day Target Level

The effect of seasonality on a fast producing yarn type (yarn type #1) with various inventory levels is shown in Figures 8, 9, and 10. There is little problem meeting demand during low demand seasons with its inventory levels at or slightly above the target NIP. During high demand seasons, there are some problems holding inventory at low target levels, however it performs well at higher levels. As with overall inventory, there is some difficulty switching from a low demand season to the high demand season as it can take anywhere up to 5 weeks to level off at the new higher inventory level

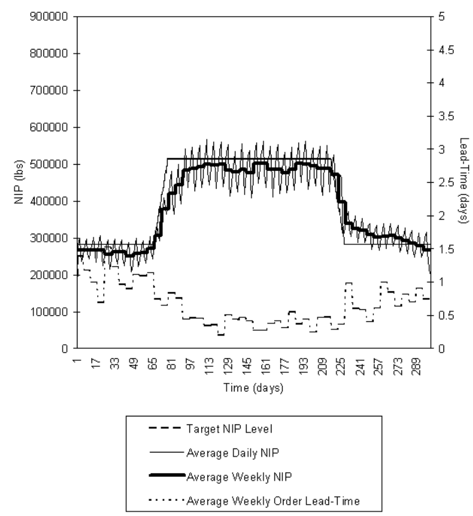

Figure 8: Seasonality for Yarn Type #1 Using a 6 Day Target Level

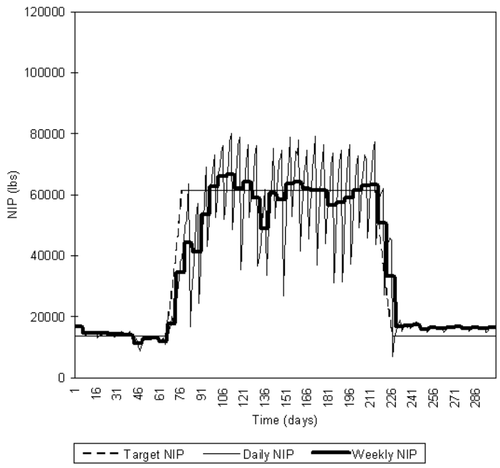

Figure 9: Seasonality for Yarn Type #1 Using a 12 Day Target Level

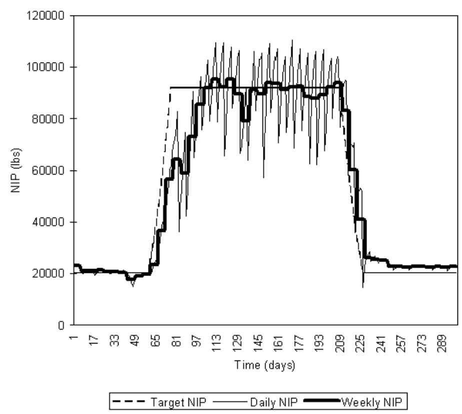

Figure 10: Seasonality for Yarn Type #1 Using a 18 Day Target Level

The seasonal effect on a slow producing yarn type (yarn type #18) is shown in Figures 11, 12 and 13. Similarly, there is no problems meeting demand for this yarn during low demand seasons. During high demand seasons, a low target level results in frequent daily stockouts while a high target level has an occasional close call.

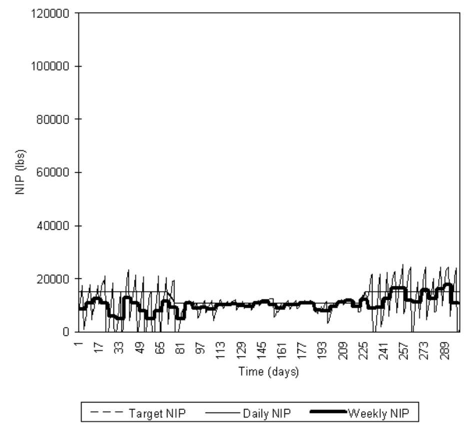

Figure 11: Seasonality for Yarn Type #18 Using a 6 Day Target Level

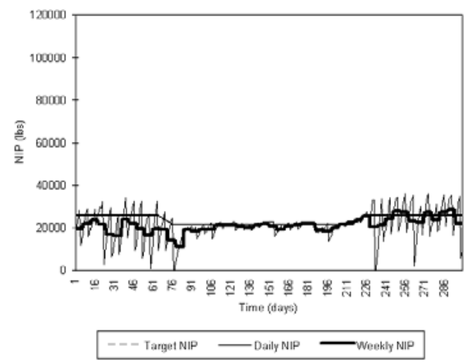

Figure 12: Seasonality for Yarn Type #18 Using a 12 Day Target Level

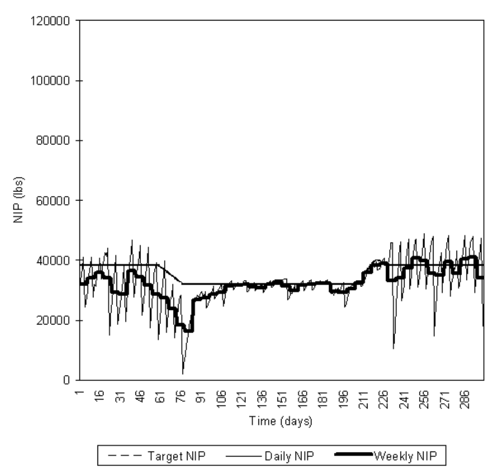

Figure 13: Seasonality for Yarn Type #18 Using a 18 Day Target Level

3.4 Conclusions

The PD algorithm reacts well to seasonal changes in demand even at small target inventory levels. Seasonality changes have little affect on order lead-time statistics. However, there is some delay in inventory levels when switching from low to high demand seasons. The switch from high to low demand seasons with regard to inventory is much quicker.

4. Conclusions and Future Research

4.1 Conclusions

Traditional min/max algorithms are often used to control yarn inventory in a textile spinning plant. Research has shown that a min/max scheme without dedicating any frames produces better on-time shipping performance than when some frames are dedicated to specific yarns. However, this comes at a cost of significantly more frame changeovers.

In this paper we develop an approach based upon the concepts of continuous process control (PD algorithm). This approach provides shipping statistics competitive with the min/max algorithm without dedicated frames and in addition has fewer frame changeovers. It determines the number of spinning frames to allocate to a yarn type based on the current and desired inventory levels through both proportional and derivative terms.

The PD algorithm would be an effective inventory control mechanism in a quick response system. It keeps inventory levels low and steady while achieving good shipping statistics and operating under a variety of plant parameters. All of these factors are important when considering a quick response system, especially for a spinning plant which is at the front of the pipeline. If a spinning plant slips on deliveries, the other plants in the pipeline must make up for the delay. The PD algorithm has been shown to provide the reliability needed for a spinning plant in a quick response system.

4.2 Future Research

To better understand how the PD algorithm would work in a quick response scenario, the model should be linked with similar models (weaving, knitting, dyeing, apparel, and retail) in the apparel pipeline. The output from one model is the input for the next model with demand rippling up from the retail customer. Only by linking all the models together can the true behavior of the quick response system be studied.

To further study the PD algorithm, a broader product mix including fashion goods should be examined. The performance of the algorithm during seasonal changeovers may be enhanced to better fit the demand.

The development of the PD algorithm is just the first step in examining the way textile plants operate. Instead of staying with tried and true inventory methods such as the min/max algorithm, plant managers should be aware other alternatives exist and consider these as tools in their re-engineering efforts to become part of a quick-response partnership.

Bibliography

King, R.E., L.A.C. Brain, and K.A. Thoney (2001), "Control of Yarn Inventory for a Cotton Spinning Plant: Part 1: Uncorrelated Demand," working paper.

Author Information

Russell E. King
Department of Industrial Engineering
Box 7906, NCSU
Raleigh, NC 27695
king@ncsu.edu

Leigh Ann C. Brain
Cary, NC

Kristin A. Thoney
Department of Textile and Apparel, Technology and Management
Box 8301, NCSU
Raleigh, NC 27695
kristin_thoney@ncsu.edu