



Control of Yarn Inventory for a Cotton Spinning Plant: Part 1: Uncorrelated Demand

Russell E. King¹
Leigh Ann C. Brain¹
Kristin A. Thoney²

¹Department of Industrial Engineering

²Department 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 this paper, a process control-based inventory control methodology is developed for a yarn spinning plant. The technique attempts to control the inventory at a specified level using separate proportional/derivative control equations for each product (yarn) type. The system is unique in that the controlling variable (number of spinning frames producing a yarn type) is discrete while the output of the control equation is continuous. Additionally, each controller is contending for a common set of resources (spinning frames). The approach is compared with a traditional min/max system via computer simulation for the case of uncorrelated demand.

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

1. Introduction

In the globally competitive textile industry the ability to meet orders in the lead-time specified, i.e., service level is critical. This is not a new phenomenon. In the late 1980's Sara Lee Knit Products (SLKP) became unhappy with their yarn suppliers. The suppliers were unable to satisfy SLKP's large demand and the average spinning machine at their mills was 20 years old. Consequently, in 1990, SLKP created Mountain City Yarn Mill so that it could reduce its dependence on external yarn sources. Spinning their own yarn

resulted in more economical and higher quality yarn (Textile World, June 1992). One approach for yarn suppliers to maintain customer satisfaction and to keep customers from spinning their own yarn like SLKP is to implement a yarn inventory system that holds the optimal amount of yarn; enough to meet demand yet small enough to be cost effective. In 1978, K. Govindarajulu and R. Nataraj studied the optimal in-process stock to be held in spinning. The in-process inventory depends upon many factors such as yarn count, type of machinery, etc.

Establishing inventory control is a necessity for reducing costs and optimizing the use of limited resources. Inventory is often the single largest asset and is a large portion of the working capital in a company. Some stock is necessary because of the highly competitive nature of the textile industry, which can result in rapid deterioration of their stocks (Bartnik, 1989). It is also necessary due to machine down time, maintenance trouble, and quality problems. Inventory can be excessive as well as inadequate. Excess inventory may result in selling pressure (Velliangiri et al., 1991). Textile mills which use modern management science tools find that the most profitable area to optimize and control is inventory (Dudeja, 1981). Some conventional techniques used in the cotton industry are stock levels and the economic order quantity (EOQ). Stock levels for safety stock, minimum and maximum stock, and reorder points avoid excess holding of inventory while minimizing stockouts. EOQ optimizes total inventory cost. The optimal level of inventory should consider forecast demand, product life, supply, price, and competitor attitude (Velliangiri et al., 1991).

Spinning plants that have routine manufacturing with the same yarns can be modeled as make-to-stock systems. With make-to-stock plants, more emphasis and control is required of production levels than in a make-to-order plant, but it is often easier to respond quickly to customer orders. The final products from the spinning plant are simply yarn of a given count and twist. Although a given product can be used to create a variety of apparel items, changes in the original demand forecast with regard to these attributes do not affect the spinning plant. However, the plant must be able to react quickly to changes involving demand in a whole product line by producing either more or less yarn of a certain count and twist. Standard techniques to determine feedback and corrective action in make-to-stock plants are forecasting, materials requirement planning (MRP) with time-phased order points, capacity requirements

planning (CRP), I/O control, work center loading, and daily dispatch reports. While these techniques are valuable in determining deficiencies and surpluses in inventory, they do not determine the allocation of resources.

To do this, a multi-product linear programming (LP) model with limited resources may be used. It can determine the amount of product to be produced in a given period. However, it assumes a known demand. Its goal is to minimize costs. The 1-product production LP model allows projected demand to differ from actual demand but allows only for one product and has no capacity limitations. For a plant to have successful production planning, a combination of these two, i.e., a multi-product, limited resources model is needed.

J
T
A
T
M
One approach is to use min/max control, i.e., maintain the level of inventory within a specified range. With the min/max algorithm, the maximum level of inventory is based on storage constraints and carrying costs, while the minimum level is chosen to control stockouts. The inventory level is checked on a user-specified frequency to see if it is above the maximum level. If so, production is stopped until the inventory falls below the minimum level.

There are two types of parameters that characterize the control. The first parameter specifies whether or not some frames are *dedicated* to high volume yarns. If so, those yarns whose average demand rate exceeds the production rate of a frame are assigned dedicated frames. For example, if the daily demand rate for some yarn is 50,000 lbs. and the daily production rate per frame is 15,000 lbs. then 3 frames would be dedicated to this yarn yielding a dedicated capacity of $3 \times 15,000 = 45,000$ lbs./day. The difference ($50,000 - 45,000 = 5,000$ in this case) would be made up on the frames not dedicated to some yarn. These frames are referred to as *scheduled frames*. The second parameter of the min/max algorithm is the target range. This refers to the difference between the maximum days of supply target that signals a stop in production of the yarn and the minimum days of supply target that signals resuming production of the yarn.

Scheduled frames are assigned yarns and operate individually. Priority on these frames is given to the yarn with the smallest net inventory position (NIP). The NIP of a yarn type is the estimated supply in inventory if the yarn is not scheduled to run on this frame. By regulating NIP, the algorithm controls the yarn in stock.

An alternative to the min/max system is to use a single target level. In this paper a proportional-derivative (PD) process control-based control algorithm is developed to keep the inventory level for each yarn as close as possible to the target using the NIP values. It is implemented using all frames as scheduled frames.

Process control is designed to control a continuously ranging system variable using a continuously ranging control variable. However, in this paper, we apply process control using multiple, discrete control variables, namely number of spinning frames running a given yarn. Since each yarn is competing for systems resources (frames), the algorithm must also handle contention issues when more frames are demanded than exist in the facility. Performance of the PD method is compared to that of the min/max algorithm in controlling the inventory of a simulated spinning cotton plant and to analyze the robustness of PD process control.

2. The Model

To compare the PD algorithm with a min/max algorithm, a stochastic simulation has been developed (Powell, 1993). It models a cotton spinning plant that produces n different yarn types. Each product is specified by a given yarn count. The plant has m spinning frames each with s spindles that are operated h hours a day, d days a week. Yarn type i is produced on a spinning frame at a rate of r_i lbs/spindle/day. Upon doffing, there is a d_c day coning/inspection delay. The yarn is then placed into inventory and is ready to ship. As previously discussed, spinning frames can be either *dedicated* or *scheduled*. A dedicated frame

is assigned a yarn type at the start of the simulation and produces only that yarn. A scheduled frame is re-assigned a yarn type periodically (every Δt time units) based on the need at that moment. A changeover delay is incurred on scheduled frames when switching from one yarn type to another. Both dedicated and scheduled frames may be idle or busy.

The rate of customer orders arriving each day is λ . Each order is randomly assigned a yarn type from a yarn mix distribution and an order quantity from an order size distribution with average order size \bar{O} . There are two types of customer orders: contract orders and spot orders. κ represents the percentage of total orders that are contract orders. Contract orders are placed by the plant's regular customers and arrive once per week. Spot orders are generated by infrequent customers and arrive at any time during the week with exponentially distributed inter-arrival times.

Orders are filled each day with a maximum of x orders filled per day to reflect shipment capacity. All remaining orders are filled at the end of the week if stock allows. No shipping is done on weekends. Contract orders have priority over spot orders, and older orders have priority over newer orders.

The plant is operated as make-to-stock where the user specifies the average safety stock in terms of days of supply D . From this, the overall desired yarn inventory level in pounds of yarn, Y , is calculated and a portion of it, y_i , is assigned to each yarn type. The goal is to satisfy customer orders within a target lead-time, as well as keep the inventory at the user specified level.

3. PD Process Control

3.1 The Algorithm

The PD algorithm is implemented every Δt time units. At time t it calculates yarn i 's NIP value, $NIP_i(t)$, and compares it with its target level, $y_{tar,i}$. The difference is called the error, e_j . By adjusting the number of frames running each yarn, the PD

algorithm attempts to alleviate the error within a user-specified amount of time. The algorithm uses process control to regulate the inventory (see Johnson 1984).

For each yarn type i , $NIP_i(t)$ is effectively a continuous function of time, including the current number of frames running yarn type i , $N_i(t)$, spindle efficiency, E , frame setup times, S , and customer order attributes, F_o , (number of customer orders, order sizes, product demand mix), i.e.,

$$NIP_i(t) = F(N_i(t), E, S, F_o).$$

Variation in these variables is responsible for creating error (actual level - desired level) in NIP. Because of uncontrollable demand variations, it is virtually impossible for the system to be error-free. Thus, the objective is to minimize the error.

The number of frames running type i , $N_i(t)$, is the controlling variable for $NIP_i(t)$ because it is the only variable that influences $NIP_i(t)$ that can also be altered by the spinner. There is a separate control equation for each yarn type. Since there is a limited number of spinning frames, contention may occur when the total number of frames required for all yarns exceeds m , the total number of frames. In this case, frames are allocated individually based on the expected time until stockout, which is a function of $NIP_i(t)$ and the average daily demand rate for yarn type.

Process control uses a repetitive procedure that involves feedback. Figure 1 is a diagram of the steps involved. The model attempts to control a continuous variable ($NIP_i(t)$) using a discrete, capacitated control variable ($N_i(t)$). Steps 1 through 4 are repeated every Δt time units and are described below.




Figure 1: Steps in the Control Process

Step 1: Measure the Controlled Variable

A raw measurement must be taken of the controlled variable. This is simply the value of $NIP_i(t)$ at the time the algorithm is being executed. $NIP_i(t)$ is defined as the current inventory level in terms of pounds of yarn, $INV_i(t)$, plus any work in process, $WIP_i(t)$ minus any outstanding orders for yarn i , $\beta_i(t)$.

$$NIP_i(t) = INV_i(t) + WIP_i(t) - \beta_i(t)$$

Step 2: Calculate the Error

$NIP_i(t)$ is compared to the target inventory level, $y_{tar,i}$, and the error is calculated.

Positive error indicates that there is too much yarn in inventory, while the error is negative when there is not enough yarn.

$$e_i(t) = NIP_i(t) - y_{tar,i}$$

Step 3: Compute the New Controlling Variable

The error must be translated into the controlling variable's units in order to correct the error. In other words, pounds of yarn must be translated into number of running frames. PD (proportional - derivative) control is used which attempts to modulate the inventory by considering the size and slope of the error curve. The control formula has a constant, proportional, and derivative term. Because the controlling variable is a discrete variable, this value is rounded.

$$c_i(t) = c_{o,i} + c_{p,i}(t) + c_{d,i}(t)$$

Constant Term, $c_{o,i}$: This term is the number of frames needed to maintain NIP_i at the target level when the error is zero. In other words, it is the number of frames needed to meet the expected daily demand and is defined as:

$$c_{o,i} = \frac{\lambda \bar{O} p_i}{sr_i E}$$

where

λ = expected number of orders/day,
 \bar{O} = expected order size,
 p_i = percentage of total demand that is for yarn type i ,
 s = number spindles per frame,
 r_i = production rate in terms of lbs./spindle day, and
 E = spindle efficiency.

Proportional Term, $c_{p,i}$: This term attempts to correct deviations from the desired level of the controlled variable in amounts directly proportional to the error, i.e.,

$$c_{p,i}(t) = K_{p,i}e_i(t)$$

where $K_{p,i}$ is a proportionality constant for yarn type i (with units of frames/lbs)

$$K_{p,i} = \frac{1}{\Delta t c_1 s r_i E}$$

The larger the absolute error, the more frames that should be allocated/deallocated. This term translates the error in terms of pounds of yarn into the equivalent number of frames required to make up that error. An input parameter, c_1 , specifies the number of Δt periods in which to make up the error. In other words, it determines how reactive the controller is to the observed error.

Derivative Term, $c_{d,i}$: The derivative term is based on the time rate of the change of the error, i.e.,

$$c_{d,i}(t) = \frac{K_{d,i}(e_i(t) - e_i(t - \Delta t))}{\Delta t}$$

where the proportionality constant

$$K_{d,i} = \frac{c_2}{s r_i E}$$

and c_2 is a number relating how important it is for the slope of the error to be zero. ($c_2 = 0$ implies having a slope of zero is not important, $c_2 = 1$ implies having a slope of zero is very important.) This term is not concerned with the actual error, but with the rate at which the error is changing, i.e., the slope of the error function at the current time. This factor attempts to stabilize the inventory level by making the slope of the error function equal to zero. Even if the

error at the present time is zero, there may still be an influence due to the momentum of the system. The logic behind the derivative term is that if the slope is increasing/decreasing there is a high probability that it will continue in that direction and therefore, fewer/more frames should be allocated to account for the observed tendency.

Integral Term: An integral term may also be included (PID control) that is based on the history of the error curve with the idea of trying to force the average error over time to zero. However, this objective is not relevant to an inventory system where customer service is an important issue and thus will not be considered further in this paper.

J T A T M Step 4: Enact the Control

Enact the new value of the controlling variable into the process by reallocating the spinning frames so that the determined number of frames, $c_i(t)$, are running yarn i . Upon completing step 3, all frames are marked for rescheduling. After each frame has doffed, a reallocation procedure that determines the yarn that the frame should produce, based on the latest values of the controlling variables, is invoked.

3.2 The Target Level

The target level, Y_{tar} , is the amount of yarn to keep in inventory to satisfy the demand for a specified number of days, D_{tar} . The average number of orders per day, λ , is multiplied by the average order size, \bar{O} , yielding the number of pounds of yarn required to meet the average demand for a day. Thus, the target level is

$$Y_{tar} = \lambda \bar{O} D_{tar}$$

The overall target level is divided among each of the individual yarns such that the probability of a stockout is the same for all yarn types (see Powell 1993).

In general, the PD controller yields average inventory levels that are higher than the target level. This is due to the fact that backorders are included in the error term.

The target algorithm seeks an inventory level that is equal to the target level assuming all backorders are filled and all *WIP* is completed. Since the shipping of orders is assumed to be spread across the entire week, yarn will be held in inventory until it is shipped. This will make the average inventory level higher than its target level.

3.3 Choosing a value for Δt

The Δt value determines how often the target algorithm is evaluated. It must be set to a value that makes sense within the context at hand. For example, assume the plant runs on a six day week where weekly contract orders arrive on Monday, the first day of the week. To react quickly to these orders, the algorithm should be called soon after these orders have arrived. Therefore, a Δt value such that the net inventory position is evaluated shortly after the arrival of the weekly contract orders is preferred.

Spot orders can arrive at anytime. The more often the target algorithm is invoked, the sooner the controller can react to these orders. Thus, in general, it is better to use smaller values of Δt .

In this experimentation Δt is set to one day. Not only does this react quickly to customer orders but it is a logical choice, i.e., re-evaluating the inventory position and determining the frame schedule for the day each morning makes more sense than re-evaluating it, say, every day and a half.

3.4 Determination of the c_1 and c_2 values

Since the system exhibits stochastic behavior, determination of the optimal parameter values (c_1 and c_2) for the algorithm is difficult. Additionally, it is desirable that the algorithm be robust, i.e., able to work well under a variety of operating conditions. Therefore, the approach taken here is to find good values for the parameters experimentally.

Generally in process control, the proportional term has the greatest impact in the ability of the controller to perform well. This term can be thought of as a macro

tuning knob while the derivative term is a fine tuning knob. The c_1 value designates the number of Δt periods in which to make up the error. The c_2 value designates how important it is to have the slope of the error curve equal to zero. This factor handles the desire to have a constant, predictable inventory with small standard deviation. In some cases, the contribution of this parameter goes against the contribution of the c_1 parameter (such as when error is negative with an increasing slope). Other times, c_2 enhances the impact of c_1 (such as when error is negative with a decreasing slope).

The Nelder Mead (Olsson 1974 and Olsson and Nelson 1975) optimization procedure can be used to find a local minima of a multi-parameter function. It is an efficient and direct search algorithm that moves towards the smaller values by moving away from high function values.

The goal is to find values for c_1 and c_2 that perform well over a large range of key input parameters. A half-factorial design was used with two levels for each of the following five factors: Number of Yarns, Capacity Utilization, Product Mix, Number of Spinning Frames, and Target Inventory Level. The function evaluated at each design point is the percent difference between the time weighted NIP value and the target level (percent error), averaged over 16 runs of length T , i.e.,

$$f = \frac{1}{16} \sum_{i=1}^{16} \frac{1}{T} \int_0^T \left| \frac{\text{NIP}_i(t) - y_{tar,i}}{y_{tar,i}} \right| dt$$

The Nelder-Mead algorithm generates a simplex of v vertices where each vertex is a function of $v-1$ variables. In this case the variables are c_1 and c_2 and the simplex has 3 vertices. Given a v^{th} order simplex, the algorithm generates another point that is a reflection of the point with the largest function value through the midpoint of the other points. The objective is to replace the point with the highest value. Based on the function evaluation at the new

point compared with those in the simplex, a new simplex is generated and the process is repeated. The algorithm is terminated when the difference between function evaluations at the simplex points is less than a given value or when the percent change within the points themselves is less than a given value.

Using the Nelder-Mead algorithm it was determined that the best point is at $c_1=3.075$, $c_2=0.00747$. Through rounding this optimal point, the simpler values of $c_1=3$, $c_2=0$ are chosen as those which will perform well under a wide range of plant configurations. Since the demand is uncorrelated, it is intuitive that the c_2 value goes to zero. If demand was correlated, then it would be expected that the c_2 value would not be zero. The correlated demand case is explored in King, Brain and Thoney, 2001.

4. Comparison of the Algorithms

4.1 Comparison Criteria

As mentioned previously, the PD algorithm tends to have an average inventory that is slightly larger than the target level. A min/max algorithm, on the other hand, tends to keep the average inventory level closer towards the minimum level. These characteristics are most obvious in cases with a large number of yarns (18-36 yarns). Because of this, the min/max algorithm cannot be accurately compared to the PD algorithm with a target level half-way between the min/max levels. In order to compare performance measures such as percentage of on-time orders and the variance of the inventory level, only cases that have approximately the same average inventory level are compared, regardless of their user-specified target or min/max levels. Separate scenarios involving 6, 12, 18, 24 and 36 yarn types were simulated over a 5-year horizon. For each scenario, the PD control algorithm was evaluated. The results were compared to min/max control where two levels of *target range* (1-day and 5-day) were tested with both dedicated and undedicated frames. In the figures that

follow, the 18-yarn scenarios are representative. For the sake of brevity, the results below are for the 18-yarn cases unless otherwise stated.

In comparing the PD and the min/max algorithms, the performance measures of interest are the percentage of on-time orders, inventory variance, and the number of frame changeovers. The first relates directly to service. Inventory variance is important because a plant that has a steady inventory level can more likely meet demand. Inventory can also be thought of as safety stock for customer orders. Therefore, lower variability in the inventory allows for lower levels of inventory to be carried without sacrificing customer service. Frame changeovers relate to the variable component of operating cost. Achieving low inventory levels with fewer frame changeovers is more economically profitable for plants.

4.2 Orders Shipped On-Time

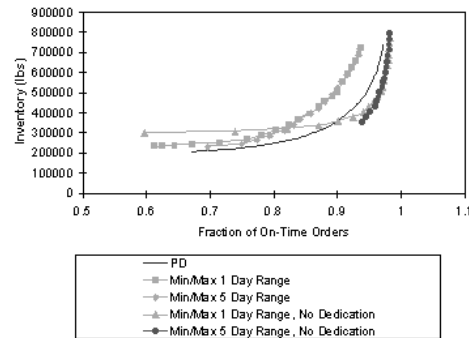


Figure 2: Percentage of On-Time Orders vs. Inventory

Figure 2 shows the average inventory levels needed for each algorithm to achieve a certain percentage of overall orders shipped within 5 days. Notice that for the min/max algorithms, the target range has little effect, but not dedicating any frames yields better on-time shipments than dedicating frames. In general, the PD algorithm is able to achieve a given level of service with less inventory than the various forms of the min/max algorithm. For

example at about 80% on-time shipping, the PD algorithm requires over 20% less inventory than any form of the min/max algorithm. However, at on-time levels above 90% the non-dedicated frame version of min/max outperforms the PD algorithm. This is achieved at the expense of frame changeovers and is discussed in section 4.4

4.3 Inventory Standard Deviation

The PD algorithm results in significantly lower inventory variance. Figure 3 shows the average standard deviation of the inventory level for each of the five algorithms tested. The PD algorithm has lower variance at any inventory level than any variation of the min/max algorithm. The standard deviations for the min/max algorithm with no dedicated frames were similar regardless of the target range. The same is true for the min/max algorithm with dedicated frames.

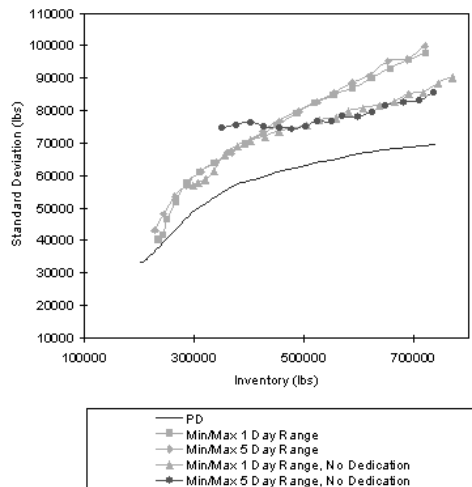


Figure 3: Inventory vs. Inventory Standard Deviation

4.4 Frame Changeovers

The no-dedicated frame versions of the min/max algorithm show better shipping performance compared to the PD algorithm at high levels of service requirement. However, as shown in Figure 4, the PD algorithm has fewer changeovers at all levels of inventory. What changeover costs are incurred by the min/max no-dedicated

frame algorithm to achieve its better shipping performance? And what does fewer changeovers mean to the PD algorithm in terms of shipping statistics? To answer these questions, Tables 1 and 2 were constructed showing 95% confidence intervals of the difference between the PD algorithm and the various min/max algorithms.

Operating with an average inventory of about 456,000 lbs., the min/max algorithm with a 1 day target range and no-dedicated frames can achieve with 95% confidence at best a 4.3% increase in on-time orders over the PD algorithm but at a cost of an extra 14.5 to 24.1 frame changeovers per day. At a high average inventory level of about 738,000 lbs, the min/max algorithm with a 5 day target range and no-dedicated frames can achieve at best a 2.09% increase in on-time orders over the PD algorithm but at a cost of an extra 11.0 to 19.1 frame changeovers per day.

J
T
A
T
M




Figure 4: Inventory vs. Frame Changeovers

Table 1: Confidence Intervals of the Differences Between the PD Algorithm and Min/Max Algorithms with Dedicated Frames

Inv. Level (1000s lbs)	Min/Max - 1 day range (Dedicated Frames)		Min/Max - 5 day range (Dedicated Frames)	
	Percent on-time orders	Daily frame changeovers	Percent on-time orders	Daily frame changeovers
236-237	9.53, 19.18	-7.3, -4.9	-	-
264-266	2.83, 11.58	-6.2, -3.8	2.18, 9.30	-5.5, -3.3
297	-	-	-	-
374-379	-	-	-	-
454-459	1.46, 9.80	-4.7, -2.3	1.54, 9.10	-4.4, -2.3
502-505	-	-	-	-
554-559	0.66, 8.41	-4.3, -2.2	0.82, 8.78	-4.1, -2.0
596-599	-	-	-	-
689-691	0.03, 7.06	-3.8, -1.8	0.17, 7.15	-3.6, -1.6
737-739	-	-	-	-

Table 2: Confidence Intervals of the Differences Between the PD Algorithm and Min/Max Algorithms with No Dedication

Inv. Level (1000s lbs)	Min/Max - 1 day range (No Dedication)		Min/Max - 5 day range (No Dedication)	
	Percent on-time orders	Daily frame changeovers	Percent on-time orders	Daily frame changeovers
236-237	-	-	-	-
264-266	-	-	-	-
297	8.29, 34.92	-31.5, -18.4	-	-
374-379	-5.61, 2.40	-26.3, -15.3	-7.44, -0.10	-24.2, -14.7
454-459	-4.32, -0.36	-24.1, -14.5	-4.85, -0.20	-23.3, -14.5
502-505	-4.15, -0.16	-22.5, -14.0	-3.9, -0.02	-22.3, -13.9
554-559	-3.59, -0.05	-21.3, -13.5	-4.24, -0.23	-20.7, -13.2
596-599	-	-	-3.10, -0.04	-20.6, -13.2
689-691	-2.18, 0.10	-19.4, -12.4	-	-
737-739	-	-	-2.09, 0.18	-19.1, -11.0

5. PD Sensitivity Analysis

While the experimental design used to test the performance of the PD algorithm encompassed input parameters such as number of frames, number of yarns, product mix and utilization, it is important to determine the impact of the individual parameters. Thus, algorithm performance is examined for robustness across a variety of input parameters. Factors considered in this section are target inventory level and the yarn to frame ratio. A plant manager must know the potential consequences on performance when considering changes in the system such as introducing new yarns or lowering the inventory level. The limits of applicability and performance of the PD algorithm are explored. An experimental design consisting of 16 runs was used for each target level between 1 and 15 days. Results for key output parameters were averaged over the 16 runs to compute a value for the given target level.

5.1 Impact of the Inventory Target Level

The results from simulations using the PD algorithm were examined to determine the impact of the target level on the performance parameters. Figure 5 shows a nearly linear relationship between the actual inventory level and the target inventory level. At smaller inventory levels, there is a slight curvature in the graph because the target level was not sufficient to meet demand. This implies that a certain level of inventory is needed to meet demand and the algorithm will produce that level even though it is higher than expected. After that level is achieved, the increase in inventory is proportional to increases in the target level of inventory as shown by the linear relationship.

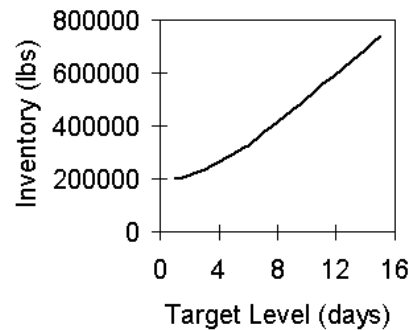


Figure 5: Target Level vs. Inventory

The absolute error between the actual level of inventory and the desired level of inventory decreases as the target level increases as shown in Figure 6. There are fewer customer backorders as inventory level increases since there is more inventory in the system to satisfy demand. Thus, the slope of the error curve begins to level off. Figure 6 shows this occurring around a target inventory of ten days. No matter what the target level, there will always be a certain amount of error in the system. This is due to the fact that demand occurs instantaneously (orders arrive as a group) while the production rate is finite. Thus, it takes some time to make the stock to return the inventory to its target level.




Figure 6: Target Level vs. NIP Error

Figure 7 shows how order lead-times are affected by target inventory level. As target inventory increases, order lead-times decrease. The largest impact is seen in the smallest inventory levels.

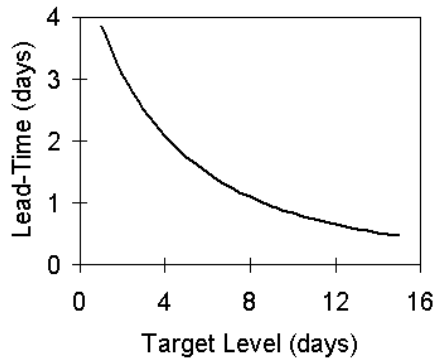


Figure 7: Target Level vs. Order Lead-Time

5.2 Impact of Yarn/Frame Ratio

Two-year simulation runs were performed varying the number of frames from 6 to 60 and the number of yarns from 6 to 36. The yarn to frame ratio has a direct impact on the on-time orders and order lead-time. Figure 8 shows an approximately linear relationship between this ratio and the order lead-time. At small ratios, the lead-time is almost variance free. As it climbs above 1, the variance increases. Thus, a plant manager can determine the number of machines needed given a target lead-time and the number of yarn types to produce. It also shows that production variety can be harmful to performance.

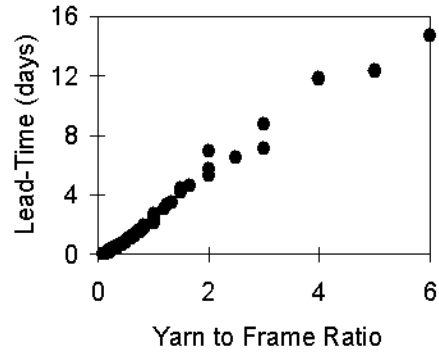


Figure 8: Yarn to Frame Ratio vs. Order Lead-Time

6. Conclusions

The PD algorithm performs much better than the min/max algorithm in terms of shipping performance, inventory variance, and frame changeovers. This algorithm results in an average inventory that is larger than the desired target level. The PD algorithm is relatively stable when considering the impact of capacity utilization, percentage of contract orders, and demand mix on overall on-time shipping performance. A target level of reasonable size (depending on the number of yarns) should be chosen to ensure a small error term as well as good shipping statistics. If the target level is too small, the percentage of on-time orders will suffer.

Much of the PD sensitivity analysis yielded expected results and thus verifies the correctness of the model. One important aspect uncovered was the importance of the yarn-to-frame ratio on performance. In the scenario used, models with a yarn-to-frame ratio under 0.75 experienced ideal performance. While the number of frames and yarns are important parameters individually, together their ratio is a determining factor of the system's performance. With this in mind, the value of this experimentation is not just in the numbers and graphs presented, but in the fact that plant managers should not necessarily focus on one parameter, but on the possible interaction of two or more parameters.

J
T
A
T
M

In part 2, a comparison of the min/max and PD-based controllers is presented under the assumption of correlated and seasonal demand processes.

Bibliography

Anon. (1992), "Mountain City: Oh, What a Yarn Mill!," *Textile World*, June, pp. 50-54.

Bartnik, Z. T. (1989), "Focus on Inventories," *Textile Asia*, February, pp 128-130.

Dudeja, V. D. (1981), "Inventory Control in Textile Mills," *The Indian Textile Journal*, August, pp. 53-59.

Govindarajulu, K., and Nataraj, R. (1978), "Optimum Process Stock in Cotton Mills," *Indian Journal of Textile Research*, 3, September, pp. 86-90

Johnson, C.D. (1984), *Microprocessor-Based Process Control*, Prentice-Hall, Inc., Englewood Cliffs, N.J.

King, R.E., L.C. Brain, and K.A. Thoney (2001), "Control of Yarn Inventory for a Spinning Plant: Part 2: Correlated Demand and Seasonality," *working paper*.

Lord, P.R. (1978), *Spinning: Conversion of Fiber to Yarn*, School of Textiles, N.C. State University.

Powell, K.A. (1993), *Interactive Decision Support Modeling for the Textile Spinning Industry*, Master of Science Thesis, Industrial Engineering Department, N.C. State University.

Velliangiri, A., and Dheenadayalu, N. (1991), "Inventory Management in Cotton Textile Industry," *The Indian Textile Journal*, December, pp. 70-76.

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

J
T
A
T
M