Int. J. Production Economics 125 (2010) 111–124

Contents lists available at ScienceDirect

Int. J. Production Economics

journal homepage: www.elsevier.com/locate/ijpe

A process control approach to tactical inventory management in production-inventory systems

Jay D. Schwartz [1], Daniel E. Rivera *

Control Systems Engineering Laboratory, School of Mechanical, Aerospace, Chemical, and Materials Engineering, Arizona State University, Tempe, AZ 85287-6106, USA

ARTICLE INFO

Article history:
Received 9 June 2008
Accepted 11 January 2010
Available online 28 January 2010

Keywords:
Supply chain management
Production-inventory systems
Production control
Process control
Internal model control
Model predictive control

ABSTRACT

Supply chain management (SCM) is concerned with the efficient movement of goods through a network of suppliers and retailers. As delayed and uncertain dynamical systems, supply chains provide an excellent opportunity for demonstrating the benefits of control engineering principles to what is traditionally perceived as a "business" problem. This paper presents a fundamental yet practical approach for applying control-theoretic principles to tactical inventory management problem in a production-inventory system, the basic unit in a supply chain. Beginning with the use of a fluid analogy, we present internal model control (IMC) and model predictive control (MPC) as means for generating a series of increasingly sophisticated decision policies for inventory management. A combined feedback-feedforward multi-degree-of-freedom IMC policy is shown to properly adjust factory starts in the presence of inventory target changes, forecasted shifts in customer demand, and stochastic changes in demand. The MPC policy displays equivalent performance, but incorporates the added functionality of managing inventory in the presence of constraints, an important practical consideration. The MPC policy shows improved performance, greater flexibility, and higher functionality relative to an advanced order-up-to policy based on control engineering principles found in the literature.

1. Introduction

As manufacturing enterprises move ahead in the 21st century, it is becoming increasingly difficult to compete on a global scale without strong inventory management strategies. The effectiveness of an otherwise innovative enterprise is limited if it is unable to deliver a desired product to customers at the correct time, to the right place, and in sufficient quantity. Therefore, production decisions must be quick, robust to uncertainty in the business environment, and optimized to meet key supply chain objectives. Improved inventory management contributes to increased revenues, lower costs, and greater customer satisfaction. There are wider societal benefits to more efficient inventory management as well. For example, minimizing the amount of unsold product reduces waste and the environmental impact of the production process, both in terms of raw materials and energy consumption. This paper presents a series of control-oriented tactical decision policies for managing inventory in production-inventory systems that ultimately seeks to advance these aims.

The study of inventory management techniques in enterprise systems has increased dramatically over the last half century (Whitin, 1970; Buffa and Miller, 1979; Seierstad and Sydsæter, 1987). The use of optimization techniques in the management of supply/demand networks began with the development of the classical Economic Order Quantity (EOQ) approaches (Wilson, 1934; Arrow et al., 1951). Since then, decision polices have been developed to accommodate dynamics (Wagner and Whitin, 1958), production costs, and inventory costs (Holt et al., 1960). Later developments include simulation-based optimization approaches for determining optimal base stock levels in "order-up-to" policies (Glasserman and Tayur, 1995; Kapuscinski and Tayur, 1998) and the application of optimal control theory (Thompson and Sethi, 1980; Sethi and Thompson, 2000) to obtain closed-form expressions for the optimal production rate. Additional control-oriented interpretations and formulations of inventory management policies have been proposed (Grubbström and Wikner, 1996; Riddalls and Bennett, 2002) with strong emphasis on mitigating demand amplification or the "bullwhip" effect (Dejonckheere et al., 2002, 2003; Towill et al., 2007; Warburton and Disney, 2007). The use of model predictive control has been popular as it offers the opportunity to robustly optimize inventories while satisfying constraints (Tzafestas et al., 1997; Perea-López et al., 2003; Braun et al., 2003; Seferlis and Giannelos, 2004; Aggelogiannaki et al., 2008). A number of excellent review articles address the application of control theory to the problem

* Corresponding author. Tel.: +1 480 965 9476; fax: +1 480 727 9321.
E-mail addresses: jay.schwartz@intel.com (J.D. Schwartz), daniel.rivera@asu.edu (D.E. Rivera).
[1] He is currently with Intel Corp., Chandler, AZ 85226, USA.

0925-5273/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.ijpe.2010.01.011

space of supply chains and production-inventory systems; these include Ortega and Lin (2004) and Sarimveis et al. (2008).

This paper is focused on how decision policies inspired from process control can be effectively applied for production-inventory and supply chain inventory management problems in uncertain, stochastic environments, as is typically the case in industrial practice. The basis for our analysis is the recognition that supply chain networks can be modeled as delayed dynamical systems. Consequently a control systems approach has significant appeal because, in contrast to classical EOQ approaches, the delayed dynamics of the system are fully incorporated in developing a decision policy. This parallels what is seen in the chemical industries, where process control systems are widely used to adjust flows to maintain level and product compositions at desired values (Ogunnaike and Ray, 1994; Seborg et al., 2004). By applying an abstraction, material flows in high-volume, discrete-parts manufacturing supply chains can be effectively modeled using a fluid analogy; one can therefore expect that decision policies based on process control principles can have a large impact on production-inventory systems and consequently, supply chain management.

A fluid representation of a production-inventory system is shown in Fig. 1. Here the manufacturing node is represented as a "pipe", while the warehouse node is represented as a "tank". Material in the pipe and tank corresponds to work-in-progress (WIP) and inventory, respectively. Specifically, in this paper we consider process control techniques such as internal model control (IMC) (Morari and Zafiriou, 1988) and model predictive control (MPC) (García et al., 1989; Camacho and Bordons, 1999) as decision policies that can provide improved performance in manufacturing systems with long throughput times and significant uncertainty, such as semiconductor manufacturing (Kempf, 2004). As control-oriented frameworks, IMC and MPC-based decision policies have the advantage that they can be tuned to provide acceptable performance in the presence of significant supply and demand variability and forecast error as well as constraints on production, inventory levels, and shipping capacity. Ultimately, the improved decision-making resulting from the application of these policies will lead to lower manufacturing costs while maximizing revenue and improving customer satisfaction.

The authors' previous work involved the development of novel model predictive control strategies for tactical decision-making in semiconductor manufacturing supply chains (Wang et al., 2007; Wang and Rivera, 2008), the optimization of these policies to meet financial objectives in a stochastic environment (Schwartz et al., 2006), and the generation of control-relevant demand forecasts tailored to these policies (Schwartz and Rivera, 2006, 2009; Schwartz et al., 2009). This paper takes a different approach. First, we focus on a standard production-inventory system and develop a series of control-oriented tactical decision policies "from the ground up," justifying for each one the distinctive functionality and unique insights brought about by applying a process control perspective. Among these include the need to apply *both* feedback and feedforward decision-making in supply chains, and the benefits of a multi-degree-of-freedom formulation for independently addressing varied supply chain requirements such as building net stock, meeting anticipated demand, and maintaining inventory in the presence of uncertain demand. The MPC decision policy is compared with a state of the art "order-up-to" policy developed by Dejonckheere et al. (2003) that is formulated using control engineering principles. By focusing on the systematic development of various control algorithms and their application on a common benchmark problem, we hope this paper will both inform the supply chain management community, and encourage greater implementation of control-oriented inventory management policies in practice.

The paper begins in Section 2 with a mathematical description of a fluid analogy for a production-inventory system. This fluid analogy forms the nominal model for the process control oriented tactical decision policies. Section 3 presents the design and application of proportional-integral-derivative (PID) control, multi-degree-of-freedom IMC, and MPC, respectively, to the control of the production-inventory system. In Section 4, the MPC decision policy is compared with the policy developed by Dejonckheere et al. (2003). Summary and conclusions are presented in Section 5.

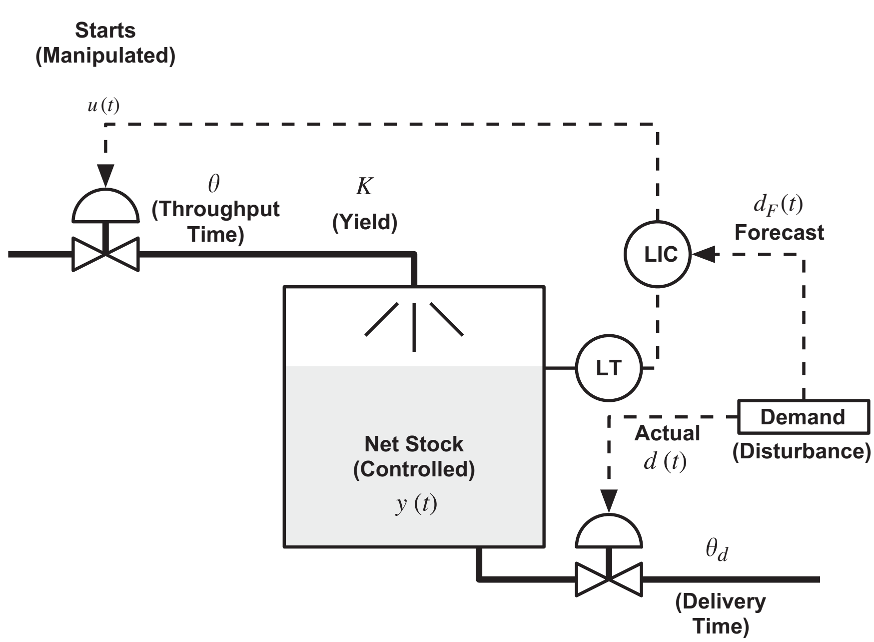

Fig. 1. Fluid analogy for a classical production-inventory system.

2. Dynamic modeling of a production-inventory system using a fluid analogy

A fluid analogy for a standard single-product production-inventory system, the simplest unit in a supply chain, is shown in Fig. 1. Fluid analogies represent meaningful descriptions of supply chains associated with high volume manufacturing problems at sufficiently long time scales, for instance, in daily or weekly decision-making. This applies to discrete-parts manufacturing problems such as semiconductor manufacturing (Braun et al., 2003). The output of a factory is stored in a warehouse where it awaits shipments to customers. The warehouse inventory serves as a buffer for stochasticity and uncertainty in both customer demand and factory production.

The factory is modeled as a pipe with a particular throughput time θ and yield K. Inventory is modeled as material (fluid) in a tank. Delivery from the warehouse is modeled as a pipe with a transportation time θ_d. Applying the principle of conservation of mass to this system leads to a differential equation relating net stock (material inventory, $y(t)$) to factory starts (input pipe flow, $u(t)$) and customer demand (output tank flow, $d(t)$) which is represented by the differential equation

$$\frac{dy}{dt} = Ku(t-\theta)-d(t) \tag{1}$$

with the corresponding transfer functions in the Laplace domain

$$y(s) = \frac{Ke^{-\theta s}}{s}u(s)-\frac{1}{s}d(s) \tag{2}$$

Net stock y is the controlled variable, while factory starts u is the manipulated variable. Based on (2) it is possible to derive feedback-only decision policies that manipulate factory starts to maintain net stock at a designated setpoint. However, if a forecast of future customer demand is available, it becomes advantageous to apply feedforward compensation. Customer demand $d(t)$ is considered as the sum of the forecasted demand ($d_F(t)$, known θ_F days ahead of time) and unforecasted demand $d_U(t)$ as shown below,

$$d(t) = d_F(t-\theta_F)+d_U(t) \tag{3}$$

Customer demand is treated as a disturbance signal at the tank outlet, which differs from traditional level control problems in the process industries where the outflow is manipulated to control the fluid level subject to disturbances at the inflow (Ogunnaike and Ray, 1994; Seborg et al., 2004). The overall dynamical system is then defined by the equation:

$$y(s) = p(s)u(s)-\overbrace{p_{d1}(s)p_{d2}(s)}^{p_d(s)}d_F(s)-p_{d2}(s)d_U(s) \tag{4}$$

$$= \frac{Ke^{-\theta s}}{s}u(s)-\frac{e^{-\theta_F s}}{s}d_F(s)-\frac{1}{s}d_U(s) \tag{5}$$

The model according to Eq. (5) is the nominal plant model for the process control-based tactical decision policies described in the ensuing sections. In these sections, a series of decision policies based on process control will be described, with each policy evaluated based on its the ability to meet inventory targets (known as setpoint tracking in the control literature) as well as satisfy forecasted and unforecasted demand changes (i.e., disturbance rejection).

3. Decision policies based on process control approaches

This section presents a variety of process-control based approaches for inventory management. It begins with a simple approach, applying a proportional-integral-derivative feedback-only policy with parameters defined using a popular tuning rule approach from the literature based on internal model control. Section 3.2 showcases a more formal (and more flexible) combined feedback-feedforward approach that is also based on internal model control. Finally, Section 3.3 discusses the basic formulation of a model predictive control-based policy. The advantages and disadvantages of each policy are discussed along with simulation results showcasing typical responses to supply chain scenarios.

3.1. IMC-based single degree-of-freedom PID control

Proportional-integral-derivative (PID) controllers are widely used in industry for a myriad of applications (Ogunnaike and Ray, 1994; Seborg et al., 2004) and are arguably the most exhaustively studied controller structure in the engineering literature. As an initial attempt, we develop and evaluate a feedback-only PID control law that manipulates the factory starts to keep the net stock at the desired target level. Fig. 2 shows a block diagram representation of such a policy. To obtain a PID form, the nominal plant model must be restricted in complexity. Using a first-order Padé approximation on the delay term leads to an integrating system with lead and lag terms,

$$p(s) = \frac{Ke^{-\theta s}}{s} \approx \frac{K\left(-\dfrac{\theta}{2}s+1\right)}{s\left(\dfrac{\theta}{2}s+1\right)} = \tilde{p}(s) \tag{6}$$

The model $\tilde{p}(s)$ shown in Eq. (6) conforms to the IMC-PID tuning rules described in Rivera et al. (1986) and Morari and Zafiriou (1988). A tutorial treatment of the application of IMC principles to PID controller design is presented in Rivera and Flores (2004). As a result of the integrating nature of the inventory process, a step change in demand becomes a Type-2 (ramp) disturbance. To guarantee offset-free control the IMC-PID tuning rule meaningful for this application must result in no offset for output ramp disturbances. Therefore, the controller must incorporate integral action, i.e. it must possess a pole at $s=0$. The IMC-PID controller corresponds to a PID with filter structure

$$u(t) = K_c\left(e(t)+\frac{1}{\tau_I}\int_0^t e(t)\,dt+\tau_D\frac{de(t)}{dt}\right)-\tau_F\frac{du}{dt} \tag{7}$$

where $e(t)$ is the error signal defined as

$$e(t) = r(t)-y(t) \tag{8}$$

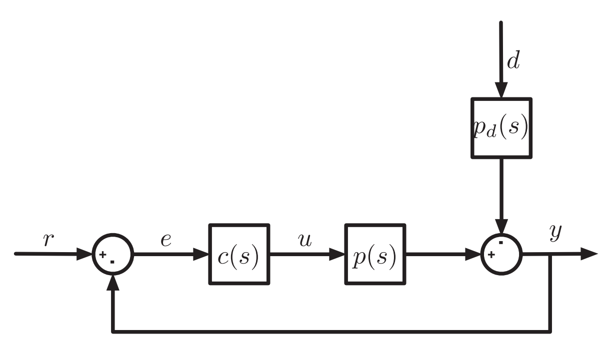

Fig. 2. Classical feedback control block diagram. The decision policy is represented by $c(s)$; $p(s)$ is the transfer function relating starts to net stock inventory, while $p_d(s)$ is the transfer function showing the effect of demand d on net stock. In a feedback control system the inventory target r is compared against the inventory signal y, and the resulting error signal e is used by the decision policy to calculate the factory starts signal u.

Eq. (7) is written in Laplace transform notation as

$$c(s) = K_c\left(1+\frac{1}{\tau_I s}+\tau_D s\right)\frac{1}{(\tau_F s+1)} \tag{9}$$

K_c, τ_I, τ_D, and τ_F are tuning parameters according to the tuning rules

$$K_c = \frac{3\theta+4\lambda}{K(\theta^2+4\theta\lambda+2\lambda^2)} \tag{10}$$

$$\tau_I = \frac{3}{2}\theta+2\lambda \tag{11}$$

$$\tau_D = \frac{\theta^2+2\theta\lambda}{3\theta+4\lambda} \tag{12}$$

$$\tau_F = \frac{\theta\lambda^2}{\theta^2+4\theta\lambda+2\lambda^2} \tag{13}$$

The tuning rule includes an adjustable parameter λ which specifies the speed-of-response and influences the performance-robustness tradeoff inherent to all closed-loop control systems. Broadly speaking, increasing λ results in a more sluggish control system response, but this decrease in performance is countered by greater robustness in the control system; that is, the control system will be stable despite mismatch between the nominal model in (6) and the true plant.

Fig. 3 showcases the performance of the IMC-PID-based tactical decision policy for changes in the inventory setpoint and demand (both forecasted and unforecasted). The production-inventory system evaluated in the simulation is characterized by a 95% process yield (K=0.95) and a production time of 3 days ($\theta=3$). The value of the user-adjustable parameter λ is 3.5 days. The simulation is comprised of an inventory setpoint change (magnitude 400) at day 20. At day 60 a customer demand step change from 0 to 100 is introduced and at day 100 the demand signal becomes stochastic. For the IMC-PID controller, the use of the Padé approximation limits the performance of the control system as there is mismatch between the internal model used to design the controller ($\tilde{p}$) and the actual plant model (p). Consequently, the controller must be detuned (increasing values of λ) to retain stability and robustness. There is significant overshoot in the net stock response to a setpoint change and the high level of "thrash" in the starts (for example, the large spike in factory starts at day 10) is unacceptable for factory managers. The presence of overshoot also drives the factory starts below zero as the controller attempts to track the target. This is a violation of physical constraints and demonstrates the need for a more accurate internal controller model and a policy with built-in constraint handling capabilities such as model predictive control. The feedback-only nature of the controller means that demand forecasts are not recognized in the decision policy, and the lack of feedforward compensation limits its ability to keep the net stock close to target levels when subjected to a significant step change in customer demand. This is reflected in the simulation shown in Fig. 3 in net stock being almost fully depleted at day 65. Better performance is expected if no approximation for delay is applied to the controller design model, and if feedforward compensation relying on demand

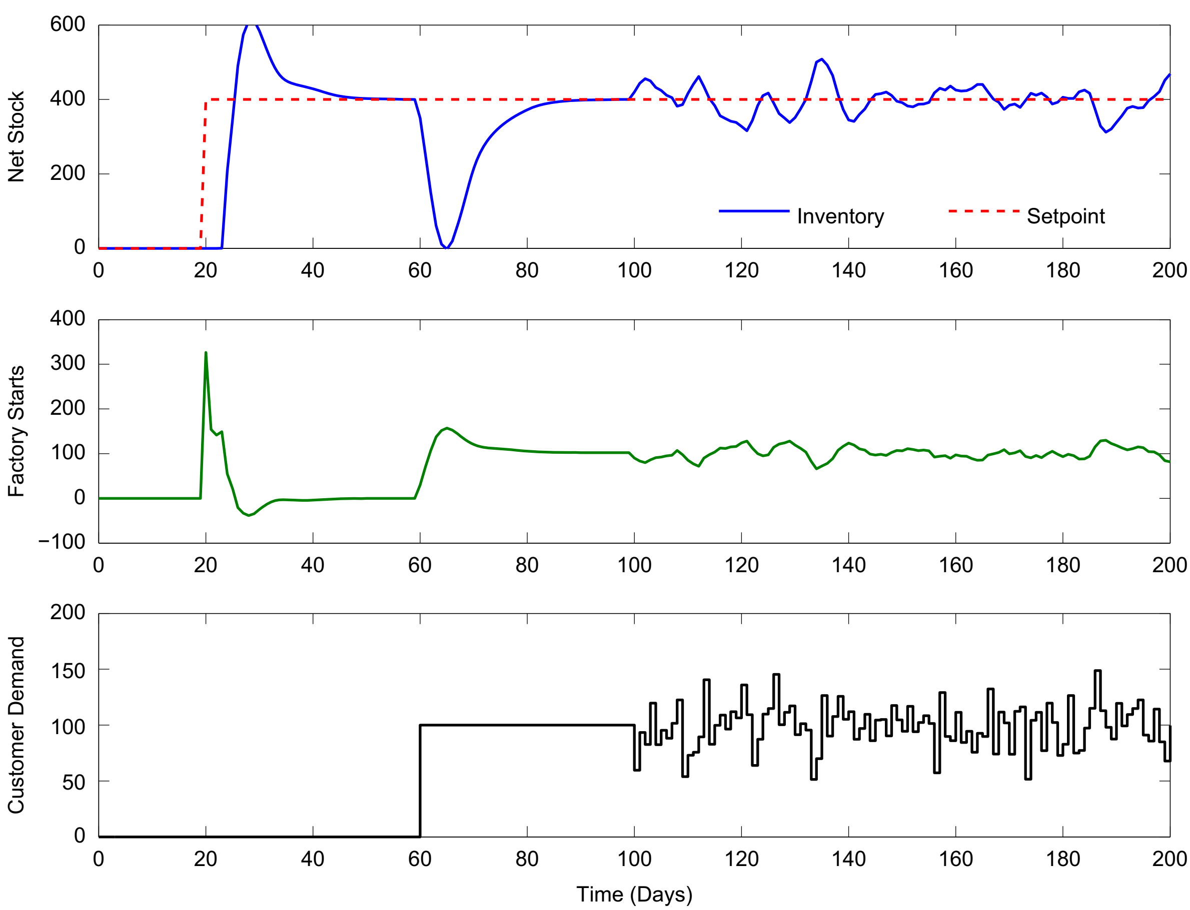

Fig. 3. One-degree-of-freedom feedback-only PID control system response. The simulation results show the inventory and starts response to a step change in inventory target at day 20, a step change in customer demand at day 60, and the introduction of uncertain, stochastic customer demand at day 100 ($\theta=3$, K = 0.95, $\lambda=3.5$).

forecasts is implemented; these issues are explored in more detail in the next subsection.

3.2. Multi-degree-of-freedom feedback-feedforward IMC

Following PID control, we consider a multi-degree-of-freedom internal model control (IMC) structure (Morari and Zafiriou, 1988) as a decision policy. With this structure independent controllers can be utilized for setpoint tracking (i.e., meeting an inventory target), measured disturbance rejection (i.e., meeting forecasted demand), and unmeasured disturbance rejection (i.e., satisfying unforecasted demand). Fig. 4 shows the structure schematically.

The nominal internal plant and disturbance models ($\tilde{p}(s)$ and $\tilde{p}_d(s)$, respectively) are integrators with delay; no approximation is used in this case:

$$\tilde{p}(s) = \frac{Ke^{-\theta s}}{s} \quad (14)$$

$$\tilde{p}_d(s) = \frac{e^{-\theta_F s}}{s} \quad (15)$$

The IMC controllers correspond to q_r for setpoint tracking, q_F for measured disturbance rejection, and q_d for unmeasured disturbance rejection which are independently tuned; hence the designation as multiple degrees-of-freedom. The IMC design procedure for these controllers is comprised the following two steps:

(1) *Design for nominal optimal performance:* $\tilde{q}_r(s)$, $\tilde{q}_d(s)$, and $\tilde{q}_F(s)$ are designed for H_2-optimal setpoint tracking, unmeasured disturbance rejection, and measured disturbance rejection, respectively. In an H_2-optimal control design the control policy is determined such that the integral square error

$$\|e\|_2^2 = \int_0^\infty e^2(t)\,dt \quad (16)$$

is minimized. Simply put, an H_2-optimal controller minimizes the average magnitude of the control error (Morari and Zafiriou, 1988). The nominal control error for the combined feedback-feedforward IMC control system when $p = \tilde{p}$ is

$$e = (1-\tilde{p}\tilde{q}_r)r - (\tilde{p}_d - \tilde{p}\tilde{q}_F)p_{d_1}p_{d_2}d_F - (1-\tilde{p}\tilde{q}_d)p_{d_2}d_U \quad (17)$$

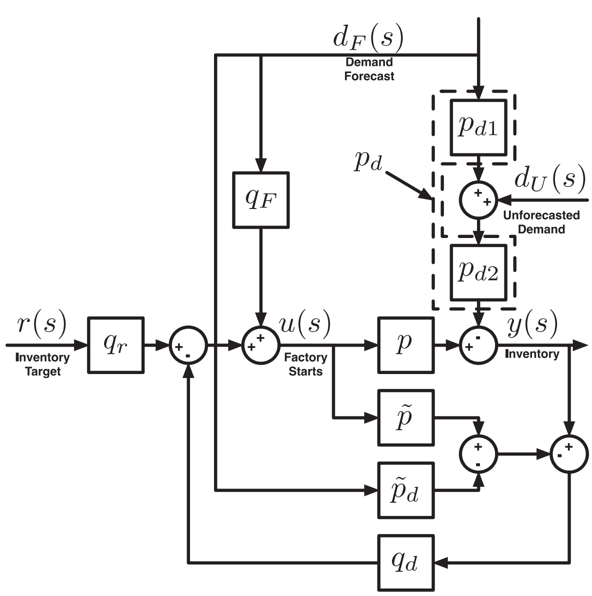

Fig. 4. Three-degree-of-freedom combined feedback-feedforward internal model control structure.

Given the flexibility afforded by the multi-degree-of-freedom formulation, it results in three distinct H_2-optimal problems that must be solved to obtain $\tilde{q}_r(s)$, $\tilde{q}_d(s)$ and $\tilde{q}_F(s)$

$$\min_{\tilde{q}_r} \|(1-\tilde{p}\tilde{q}_r)r\|_2 \quad (18)$$

$$\min_{\tilde{q}_d} \|(1-\tilde{p}\tilde{q}_d)p_{d_2}d_U)\|_2 \quad (19)$$

$$\min_{\tilde{q}_F} \|(\tilde{p}_d-\tilde{p}\tilde{q}_F)p_{d_1}p_{d_2}d_F\|_2 \quad (20)$$

subject to the requirement that $\tilde{q}_r(s)$, $\tilde{q}_d(s)$ and $\tilde{q}_F(s)$ be stable and causal. For problems (18)–(20) to be well-posed, specific forms for the setpoint/disturbance signals r, d_U, and d_F need to be specified.

(2) *Design for robust stability and performance:* In this step $\tilde{q}_r(s)$, $\tilde{q}_d(s)$ and $\tilde{q}_F(s)$ obtained from (18) to (20) are augmented with low-pass filters which can be tuned to detune the nominal performance (e.g., reduce aggressive manipulated variable action associated with the optimal controller per Step 1) or to satisfy robust performance.

The final controllers obtained from applying this procedure to the models according to (14) and (15) are shown as follows:

Setpoint tracking: The setpoint tracking mode of this control system is designed using an H_2-optimal controller for a step change, augmented with a first-order filter. The controller guarantees no offset for Type-1 (step) setpoint changes in the control system. The mode allows the controller to adjust safety stock inventory targets to any user-desired level.

$$q_r(s) = \frac{s}{K}\frac{1}{(\lambda_r s+1)^{n_r}} \quad (21)$$

Unmeasured disturbance rejection: This mode of the control system allows the user to specify the system response to unforecasted demand changes. The design procedure relies on an H_2-optimal controller for *ramp* disturbance changes, with a generalized Type-2 filter guaranteeing no offset for both asymptotically step and ramp disturbances,

$$q_d(s) = \frac{s(\theta s+1)}{K}\frac{(n_d\lambda_d s+1)}{(\lambda_d s+1)^{n_d}} \quad (22)$$

Measured disturbance rejection: The measured disturbance rejection mode of the control system performs feedforward control action by relying on a θ_F- day ahead forecast signal to manipulate factory starts. The IMC controller form is defined as follows (Lewin and Scali, 1988):

$$q_F(s) = q'_F(s)f_F(s) \quad (23)$$

where $q'_F(s)$ consists of

$$q'_F(s) = (e^{-(\theta_F-\theta_d-\theta)s}/K) \quad (24)$$

when the forecast horizon is longer than the sum of the factory throughput time and delivery time ($\theta_F \ge (\theta+\theta_d)$). If the forecast horizon is shorter ($\theta_F \le (\theta+\theta_d)$) then $q'_F(s)$ consists of

$$q'_F(s) = ((\theta+\theta_d-\theta_F)s+1)/K \quad (25)$$

The generalized Type-2 filter $f_F(s)$ is defined as

$$f_F(s) = \frac{(n_F\lambda_F s+1)}{(\lambda_F s+1)^{n_F}} \quad (26)$$

Each controller is required to be stable and proper, thus imposing the restriction that all values of the user adjustable parameter λ be positive ($\lambda_i > 0$) and that the filter order is chosen to ensure transfer function properness ($n_r \ge 1$, $n_d \ge 3$, $n_F \ge 3$). Normally the filter design is restricted to semiproper forms for q_r and q_d and to the strictly proper case for q_F.

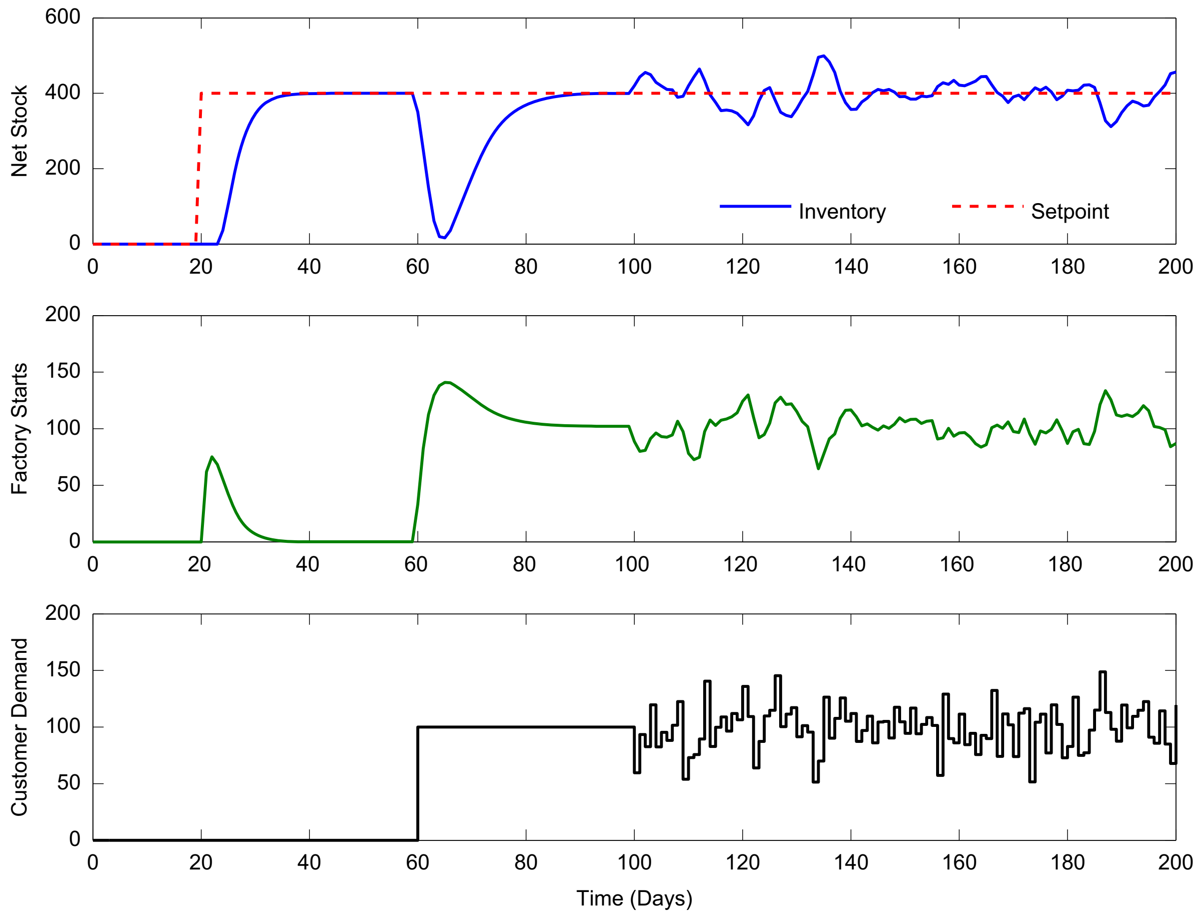

Fig. 5. Two-degree-of-freedom feedback-only IMC response to an inventory setpoint change, a forecasted demand change, and unforecasted stochastic customer demand ($\theta = 3$, $\theta_F = 5$, $K = 0.95$, $\lambda_r = 2$, $n_r = 2$, $\lambda_d = 4$, $n_d = 3$).

The responses of the multi-degree-of-freedom policies for the simulation parameters considered in Section 3.1 are shown in Figs. 5 and 6. Fig. 5 shows the response of a two degree-of-freedom (2DoF) feedback-only control system that does not take advantage of a demand forecasts, while Fig. 6 considers the fully functional, three degree-of-freedom (3DoF) combined feedback-feedforward control system. Because these controllers were designed using the actual plant model as opposed to the Padé approximated model, the closed-loop speed-of-response can be increased substantially without encountering stability problems. The added degrees-of-freedom in these tactical decision polices allow for independent tuning of the setpoint tracking and disturbance rejection responses. As a result, these multi-degree-of-freedom policies do not exhibit overshoot when subjected to setpoint changes. Both the 2DoF and 3DoF control systems perform equivalently on unforecasted demand changes; however, their performance varies greatly in the presence of forecasted demand. The combined feedback-feedforward 3DoF policy anticipates the demand change at day 60 and adjusts starts accordingly, while the feedback-only 2DoF policy reacts to the demand change once it occurs; given the three days of throughput time this results in a need for high levels of safety stock. The substantially superior performance resulting from feedforward control action highlights the need for the effective use of demand forecasts in managing supply chains. By applying feedforward compensation, inventory targets can be lowered substantially, and starts changes in response to demand changes can be less abrupt, consequently decreasing inventory holding costs, smoothing factory operations, and improving profitability.

While the application of IMC-based algorithms provides insight into the design of robust, control-oriented policies for inventory control in a production-inventory system, ultimately our interest is in discrete-time (i.e., sampled-data) controllers capable of handling constraints and being scaled to larger networks. The next subsection presents a model predictive control strategy that meets these requirements.

3.3. *MPC as a tactical decision policy*

Model predictive control (MPC) stands for a family of methods that select control actions based on on-line optimization of an objective function. MPC has gained wide acceptance in the chemical and other process industries as the basis for advanced multivariable control schemes (García et al., 1989; Camacho and Bordons, 1999). In MPC, a system model and current and historical measurements of the process are used to predict the system behavior at future time instants. A control-relevant objective function is then optimized to calculate a sequence of future control moves that must satisfy system constraints. The first predicted control move is implemented and at the next sampling time the calculations are repeated using updated system states; this is referred to as a moving or receding horizon strategy. Fig. 7 presents a useful visualization of the MPC approach. A demand forecast signal is used in the moving horizon calculation to anticipate future system behavior, which plays a significant role in the use of MPC for supply chain applications. The appeal of MPC over traditional approaches to feedback and feedforward control

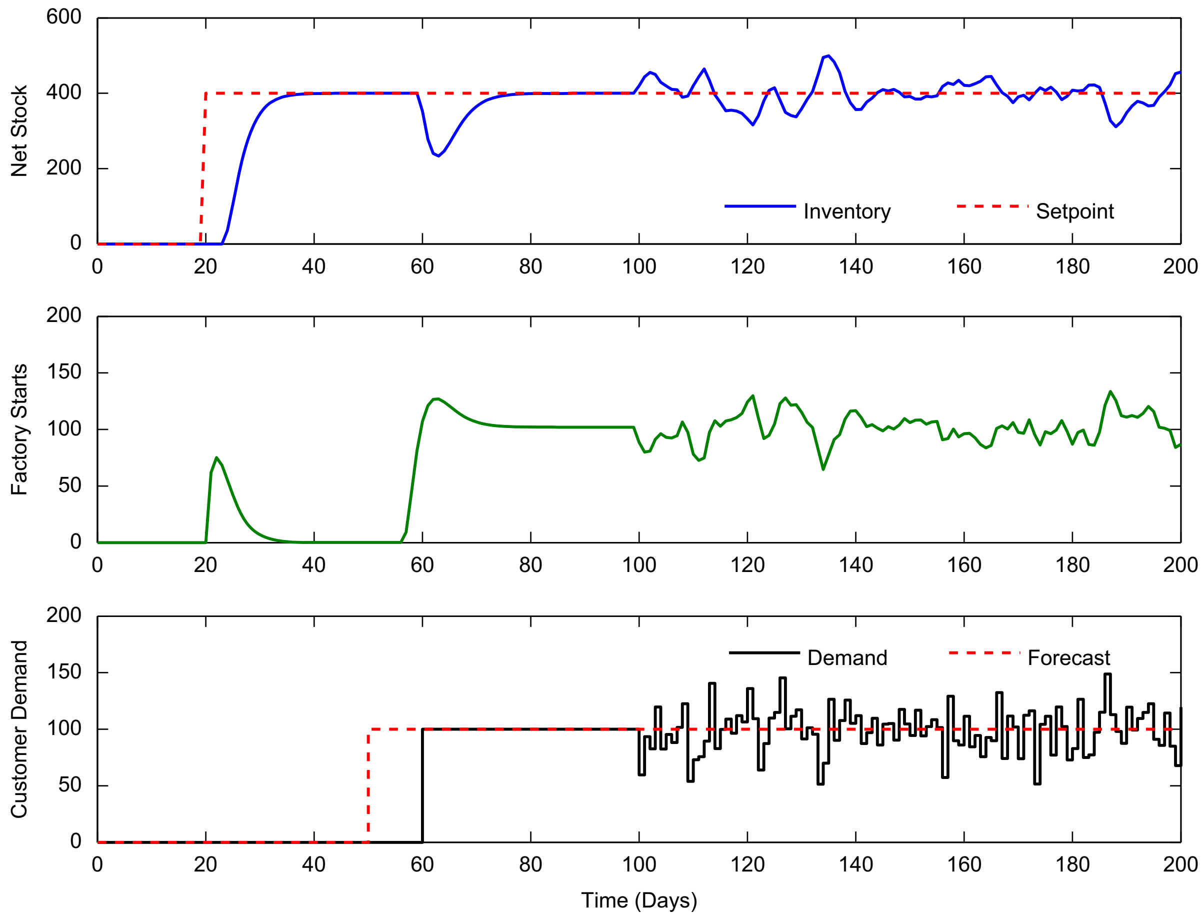

Fig. 6. Three-degree-of-freedom combined feedback-feedforward IMC response to an inventory setpoint change, a forecasted demand change, and unforecasted stochastic customer demand ($\theta=3$, $\theta_F=5$, K = 0.95, $\lambda_r=2$, n_r = 2, $\lambda_F=2$, n_F = 3, $\lambda_d=4$, n_d = 3).

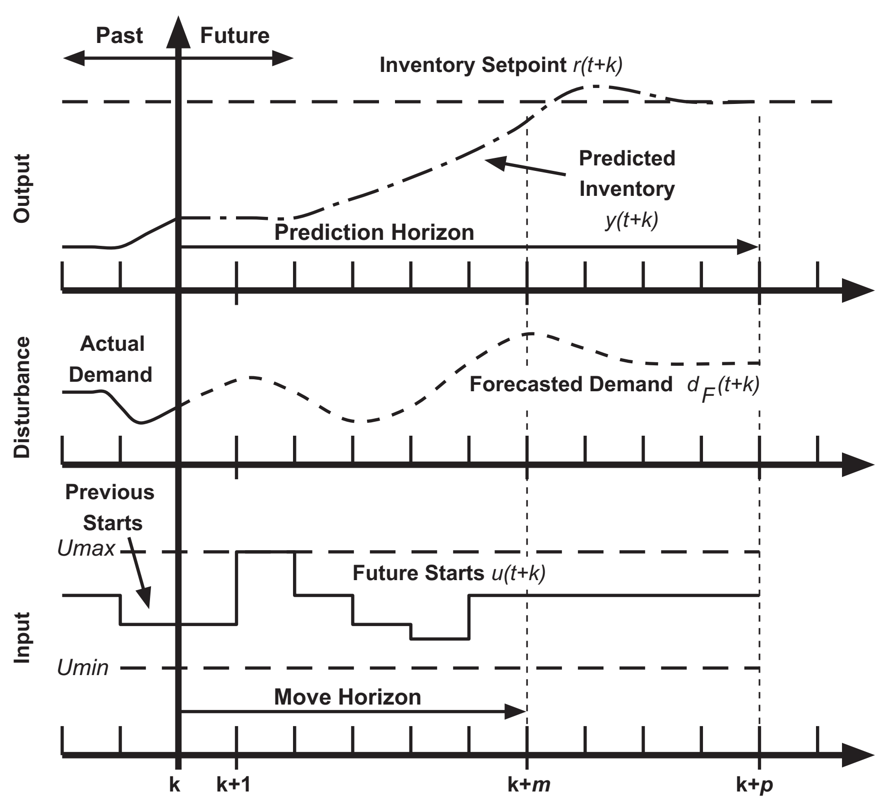

Fig. 7. Receding horizon diagram for model predictive control.

design includes the explicit handling of constraints on system input and output variables and its relative ease-of-use. In this latter category, MPC represents a completely "time-domain" technology that can be readily used in multivariable problem settings and has received wide appeal in the process industries (Qin and Badgwell, 2003). Whereas the IMC policies presented in Sections 3.1 and 3.2 were presented in continuous-time, the nominal process model for an MPC policy is expressed in sampled (i.e., discrete) time, and corresponds to

$$y(k+1)=y(k)+Ku(k-\theta)-d_F(k-\theta_F)-d_U(k-1) \tag{27}$$

k is an integer representing the sampling instant for a sampling time T; for simplicity we will assume that θ and θ_F are integer multiples of T.

The flexibility to define objective functions is an important aspect of MPC; a meaningful formulation for the inventory management problem considered in this paper is as follows:

$$\min_{\Delta u(k|k)\cdots\Delta u(k+m-1|k)} J \tag{28}$$

where the individual terms of J correspond to

$$J=\sum_{\ell=1}^{p} Q_e(\ell)(\hat{y}(k+\ell|k)-r(k+\ell))^2+\sum_{\ell=1}^{m} Q_{\Delta u}(\ell)(\Delta u(k+\ell-1|k))^2 + \sum_{\ell=1}^{m} Q_u(\ell)(u(k+\ell-1|k)-u_{target}(k+\ell-1|k))^2 \tag{29}$$

subject to constraints on inventory capacity,

$$0 \le y(k) \le y_{max} \tag{30}$$

starts inflow capacity,

$$0 \le u(k) \le u_{max} \tag{31}$$

and changes (i.e., "moves") in the number of factory starts per sampling instant,

$$\Delta u_{min} \le \Delta u_k \le \Delta u_{max} \tag{32}$$

Here p is the prediction horizon and m is the control horizon. Q_e, $Q_{\Delta u}$, and Q_u are penalty weights on the control error, move size and control signal, respectively. The MPC algorithm can be enhanced with WIP as an additional output variable subject to constraints, as shown in Wang et al. (2007).

Eq. (29) is a multi-objective expression that addresses the main operational objectives in the supply chain. The first term is a setpoint tracking term that is intended to maintain inventory levels at user-specified targets over time. The second term is a move suppression term that penalizes changes in the starts. Penalizing changes in the starts rate, while desirable from the standpoint of factory operations, also serves an important control-theoretic purpose, as a mechanism for introducing robustness in the control system (García et al., 1989). The third term is an input target term that is intended to maintain the starts response close to daily target values, based on targets calculated at the strategic level. The emphasis given to each one of the sub-objectives in (29) (or to specific system variables within these objective terms) is achieved through the choice of weights ($Q_e(\ell)$, $Q_{\Delta u}(\ell)$, and $Q_u(\ell)$).

The solution to (29) with a linear model such as (7) and inequality constraints according to (30)–(32) can be solved using quadratic programming (QP) (Camacho and Bordons, 1999). A flowchart describing the sequence of steps involved in an MPC control calculation is depicted in Fig. 8. At each sampling instant the MPC controller obtains a fresh net stock measurement, then relies on the model along with knowledge of prior measurements and controller moves to make a prediction of the inventory level into the future for a horizon length p. The difference between this prediction and a desired reference level, referred to as the error projection, is provided to the QP algorithm. The solution to a quadratic program yields the optimal move sequence for a horizon m that minimizes the objective function, given a set of user-specified weights $Q_e(\ell)$, $Q_{\Delta u}(\ell)$, and $Q_u(\ell)$. The first move in the sequence is then implemented as the starts value for that time interval, with the process being repeated at the next sampling instant in a receding horizon fashion.

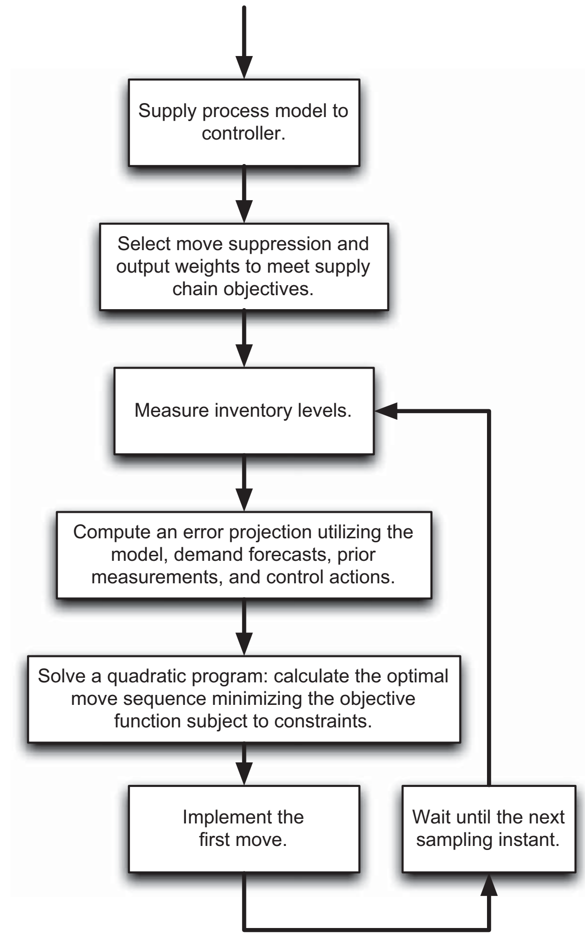

Fig. 8. Flow diagram for the MPC decision policy. Once a process model and tuning have been established the controller will measure inventory levels and use the process model to make a prediction of future inventory levels. The factory starts are then optimized over a horizon to meet objectives defined by the tuning. The first move is implemented and the process is repeated.

The MPC formulation presented in this paper follows the classical MPC problem solution (García et al., 1989; Camacho and Bordons, 1999), but is different in that it utilizes anticipation of disturbance signals. A multi-degree-of-freedom formulation patterned after the one presented for IMC has been developed (Wang, 2006; Wang and Rivera, 2008) but is not considered in this paper. In the presented implementation the model used for prediction is fixed throughout a simulation; this can be contrasted with an online production process identification approach (Aggelogiannaki et al., 2008). All simulations shown in this paper utilize the MATLAB® Model Predictive Control Toolbox. Model predictive control has been successfully applied to multi-echelon supply chains by the authors and other researchers (Braun et al., 2003; Perea-López et al., 2003; Seferlis and Giannelos, 2004; Schwartz et al., 2006; Wang et al., 2007; Wang and Rivera, 2008).

Fig. 9 shows the response of an MPC-based decision policy that is contrasted to the combined feedback-feedforward IMC design of Fig. 6. Factory starts begin to increase at day 20 in response to

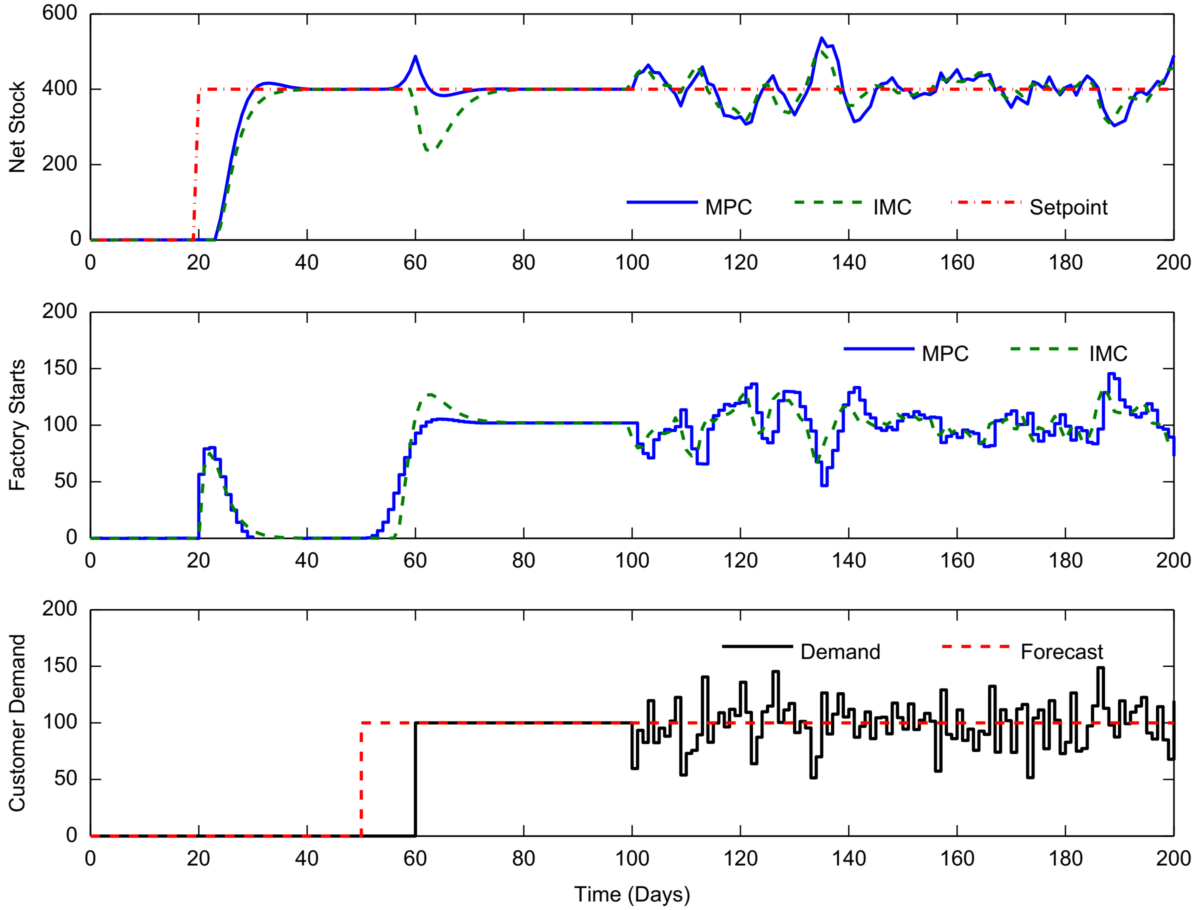

Fig. 9. Solid lines: MPC-based response with demand change anticipation (Q_e = 1, $Q_{\Delta u}=5$, Q_u = 0, p = 10, m = 5). Inventory level setpoint is 400 units. Dashed-lines: combined feedback-feedforward three-degree-of-freedom IMC-based policy response ($\lambda_r=2$, n_r = 2, $\lambda_F=2$, n_F = 3, $\lambda_d=4$, n_d = 3).

a net stock setpoint change; both IMC and MPC perform equivalently in setpoint tracking. Differences between the policies are best seen for the case of forecasted demand changes. For the MPC policy, factory starts increase at day 52 in anticipation of a demand change occurring at day 60. In contrast to the IMC-based policy, the MPC-objective function simultaneously penalizes both inventory deviation from setpoint and changes (i.e., moves) in factory starts. As shown in Fig. 9, the MPC decision policy builds inventory in anticipation of the demand change while the IMC-based approach adjusts starts exactly θ days before the demand change occurs, depleting existing net stock in the process. One could justify the merits of each of these differing responses, with the MPC approach possessing the advantage that by immediately building inventory to meet a future demand it decreases the possibility of stockouts. Unforecasted stochastic customer demand is introduced at day 100, and here both the IMC and MPC decision policies display equivalent performance.

3.4. MPC response under constrained conditions

Fig. 10 contrasts the MPC response shown in Fig. 9 with an identically tuned MPC policy under constrained conditions. The constrained MPC is subject to a maximum factory starts level constraint of 110, a maximum rate of change starts constraint of 10/day, and a maximum inventory limit of 465 units. In the simulation of Fig. 10 an anticipated step setpoint change occurring at day 20 is evaluated, followed by a forecasted demand change occurring at day 60. The constrained MPC policy is able to consistently track the inventory target throughout the simulation while also minimizing changes to the factory inputs. Enforcing the maximum inventory limit creates an unavoidable need to dig deeper into safety stocks. The built-in constraint handling of MPC represents a practical advantage for use of these algorithms in real-life manufacturing problems.

4. MPC comparison with the smoothing replenishment rule

In this section we compare the MPC-based decision policy and an advanced order up-to policy developed by Dejonckheere et al. (2003), henceforth referred to as the smoothing replenishment rule (SRR). As noted in Dejonckheere et al. (2003), the formulation of SRR relies on control engineering principles, with a block diagram for the policy shown in Fig. 11. T_p is the production time, while T_n is a user-adjustable parameter that dictates the response to a discrepancy of finished goods net stock. T_w is a similar parameter that defines the response to an on-order position discrepancy. The parameter α determines the bandwidth of an exponential smoothing filter applied to the actual demand signal; this constitutes the forecasting policy for the algorithm. The SRR policy lacks a setpoint tracking mode that would enable it to accept user-defined inventory targets, and relies on a net stock target that is determined from the smoothed demand signal. Exponential smoothing also causes the policy to lag substantially when a step demand change is introduced. Fig. 12 compares these performance characteristics of the SRR policy (using parameters T_a = 8, T_n = 4, and T_w = 4 as defined in Dejonckheere et al., 2003)

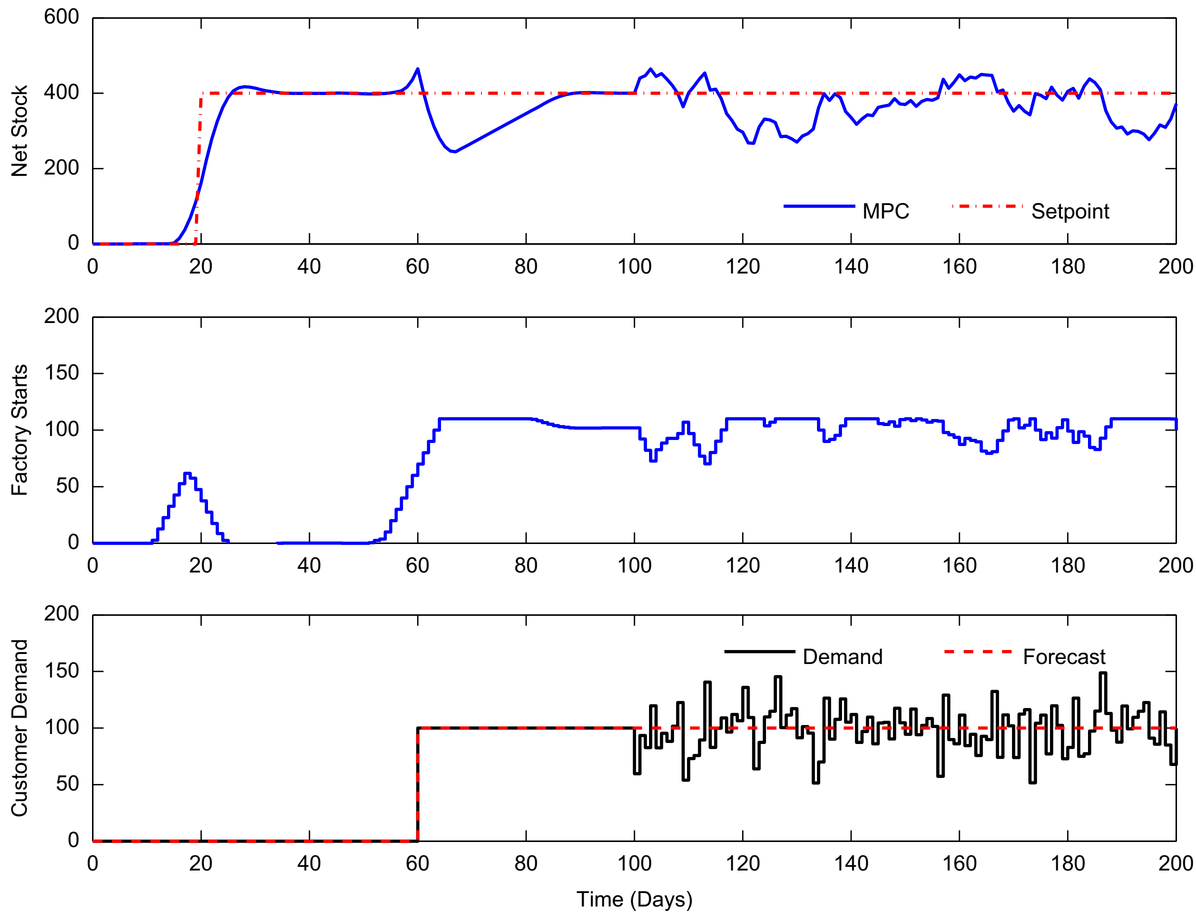

Fig. 10. Solid lines: MPC-based response with demand and setpoint change anticipation (Q_e = 1, $Q_{\Delta u} = 5$, Q_u = 0, p = 10, m = 5) operating under constraints: maximum starts level of 110 units ($u \leq 110$), maximum starts change rate of 10 units/day ($|\Delta u| \leq 10$): and a maximum inventory value of 465 units ($y \leq 465$).

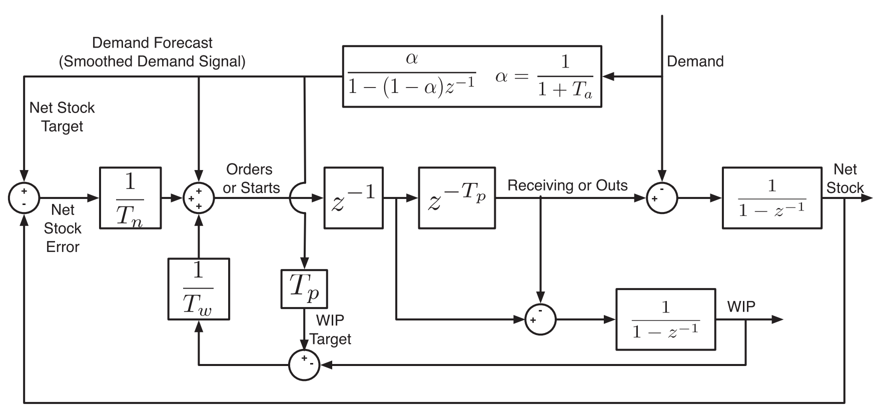

Fig. 11. Block diagram for the smoothing replenishment rule (SRR) as developed by Dejonckheere et al. (2003).

with those of MPC, the latter taking full advantage of demand forecasts through anticipation in the receding horizon algorithm.

The SRR and MPC policies possess equivalent functionality in the case of stationary unforecasted customer demand; Fig. 13 shows simulation results comparing the performance of the two policies under these conditions. While both policies are able to manage the inventory level without significant factory thrash, the MPC policy results in fewer stockouts and a smoother factory starts profile. These results motivate a more exhaustive study of the policies for a wider range of tuning parameters; this is summarized in Fig. 14. Inventory deviation from setpoint is plotted along the vertical axis and factory starts change variability (factory thrash) is plotted along the horizontal axis. An ideal inventory management policy would display low levels of inventory deviation and factory thrash, resulting in values close to the lower left hand corner of the plot. Each data point on the plot represents the average of 100 simulation runs equivalent to those shown in Fig. 13. In all simulations, the throughput time (T_p, θ) and yield were specified as 3 days and K = 1, respectively. The results for the smoothing replenishment rule (red circles) are shown for all tunings in the following ranges: $6 \leq T_a \leq 10$, $2 \leq T_n \leq 6$, and $2 \leq T_w \leq 6$. Move suppression values $Q_{\Delta u}$ for the MPC policy vary from 1 to 50; all other parameters are kept at fixed values (output weight Q_e=1, prediction horizon P=10 days,

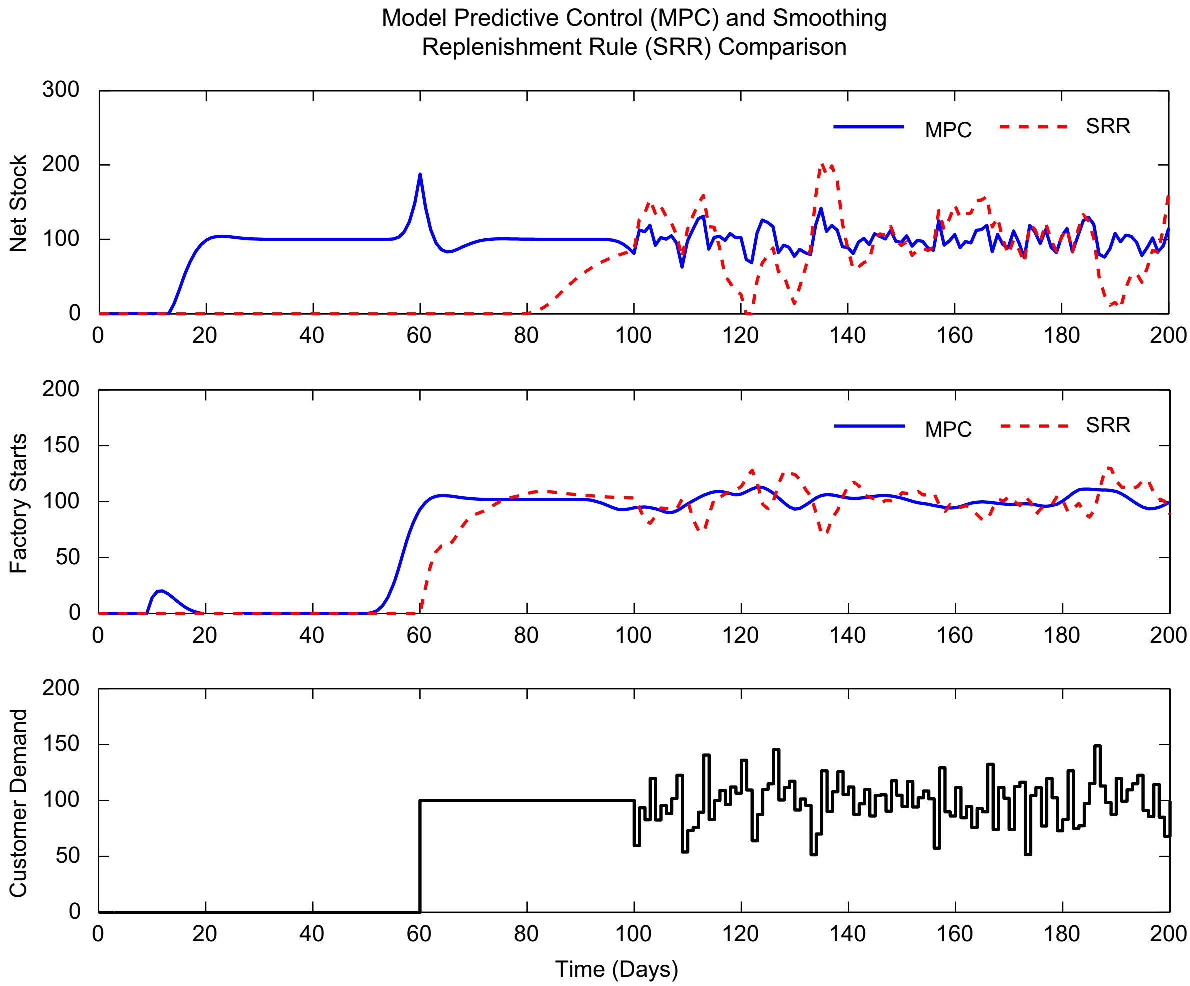

Fig. 12. Testing the MPC and SRR decision policies: step setpoint and demand changes, and stochastic unforecasted demand. Both MPC and SRR are able to manage inventory in the presence of stochastic customer demand. However, the MPC policy is capable of tracking a user-defined setpoint and takes full advantage of anticipated demand forecasts. This simulation was performed on a production-inventory system with a production time of 3 days and K = 1. MPC parameters are as follows: move suppression $Q_{\Delta u}=5$, output weight Q_e = 1, prediction horizon p = 10, and move horizon m = 5. Parameters for the SRR are T_a = 8, T_n = 4, and T_w = 4 as defined in Dejonckheere et al. (2003).

and move horizon M=5 days). All policies were evaluated using the demand signal

$$d(k)=1000+25\frac{1}{1-0.9q^{-1}}a_1(k)+25a_2(k) \quad (33)$$

where a_1 and a_2 are normally distributed white noise sequences, and q^{-1} is the backshift operator. Two MPC policies, with and without demand anticipation, were studied. For the MPC policy with anticipation the following demand forecast signal was used:

$$d_F(k)=1000+25\frac{1}{1-0.9q^{-1}}a_1(k) \quad (34)$$

where a_1 is the same realization of the white noise sequence used to calculate the actual demand.

Examining Fig. 14 shows that the MPC policy has two principal advantages over SRR: improved performance over a wider range of tuning parameters, and more intuitive, flexible tuning. Consider the latter criterion. The Smoothing Replenishment Rule has three adjustable parameters that impact performance. MPC tuning can be reduced to a single parameter, the move suppression weight $Q_{\Delta u}$, which can be used to tradeoff inventory deviation from setpoint with factory starts change variability. The MPC performance curve (blue crosses) shows the greater flexibility of the MPC policy. A planner can specify MPC to operate anywhere along this curve by adjusting the move suppression value, whereas the SRR policy is restricted to a narrower performance region, indicated by the red circles. For a move suppression value of 20 and without anticipation the MPC policy results in an inventory deviation from target value of 4463 and a factory starts change variability value of 269. The smoothing replenishment rule can be tuned to provide comparable performance to the MPC operating under these conditions. Using the SRR tuning parameters discussed in Dejonckheere et al. (2003) and Fig. 13 gives 5041 for the inventory deviation metric and 344 for the factory thrash metric. The use of anticipation in conjunction with the MPC policy substantially improves supply chain performance, as shown by the downward and leftward shift of the performance curve (green squares). Despite the presence of forecast error the MPC policy with anticipation results in both less inventory deviation and factory thrash. For a move suppression value of 20, using anticipation causes the inventory deviation metric to decrease 56.3% to 1949 and the factory thrash metric to decrease 44.6% to 149.

5. Summary and conclusions

Improved management of production-inventory systems has been the topic of significant study for several decades, with

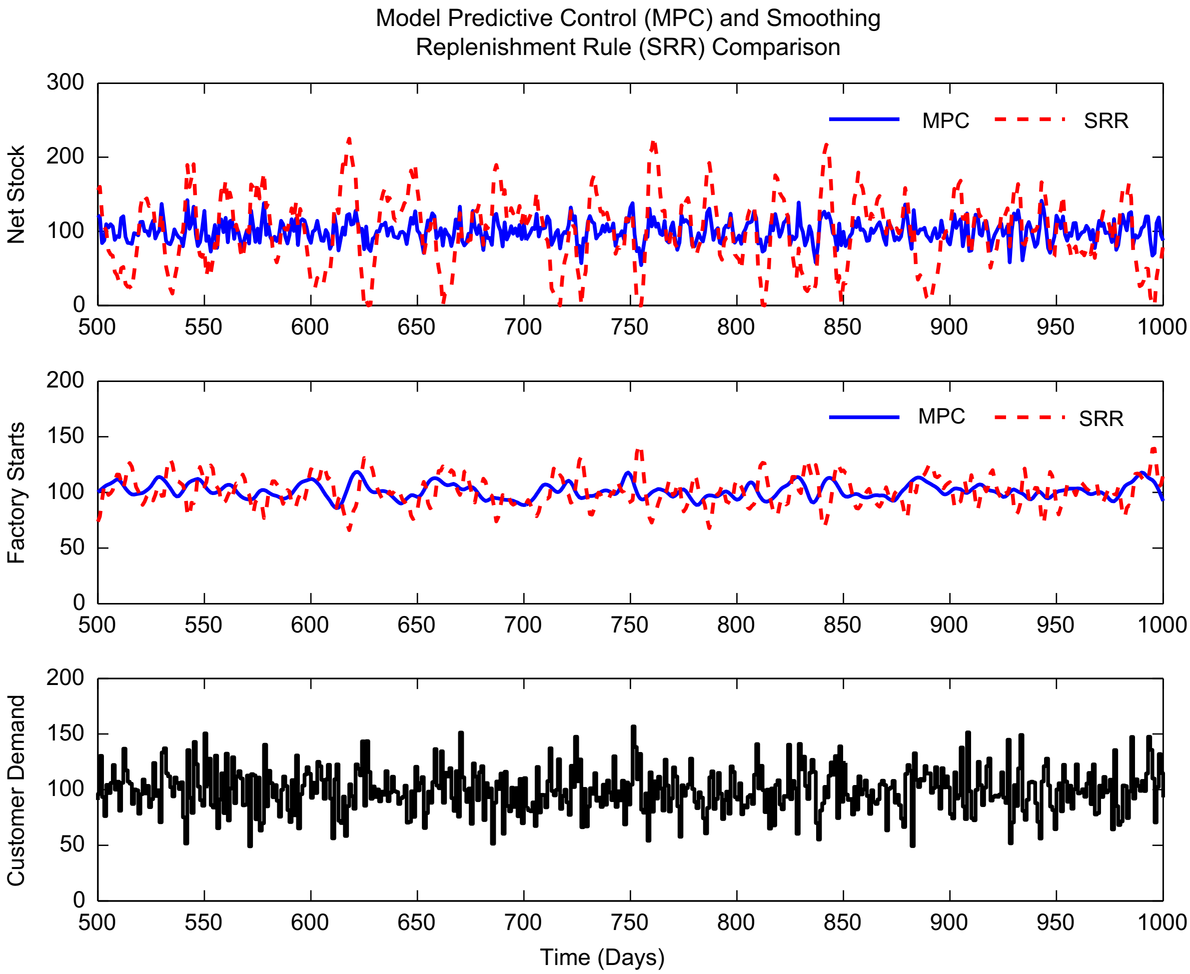

Fig. 13. Decision policy comparison for MPC and SRR under conditions of stochastic, uncertain customer demand. The MPC policy provides closer inventory target tracking with less factory thrash. System and tuning parameters are equivalent to those used in Fig. 12.

control engineering approaches forming part of the plethora of methods that have been applied to the problem. To adequately manage inventory in the presence of uncertain customer demand and stochastic conditions it is important for a decision policy to properly capture the underlying dynamics of the production-inventory system. Applying a fluid analogy to the manufacturing process yields well-understood mathematical expressions for the system dynamics that serve as nominal models for tactical inventory management policies based on process control principles. A series of control-oriented decision policies of increasing sophistication and functionality were developed in this paper, based on the concepts of internal model control and model predictive control. Their effectiveness was evaluated in simulations involving changes in net stock targets (setpoint tracking), response to forecasted demand (measured disturbance rejection) and response to unforecasted demand (unmeasured disturbance rejection).

A feedback-only PID strategy with IMC tuning, while stable, suffered performance limitations arising from relying on a nominal process model of restricted complexity. A feedback-only, two-degree-of-freedom IMC policy improved the response to setpoint tracking, and enabled the setpoint tracking and unforecasted demand responses to be tuned independently. However, lacking the ability to incorporate a demand forecast, the feedback-only policies require large safety stocks to avoid running out of product during significant customer demand increases. These limitations were addressed by implementing a combined feedback/feedforward, three-degree-of-freedom controller that makes full use of demand forecasts. The IMC formulation provides valuable insights into inventory control problems; however, IMC policies may perform poorly in the presence of constraints, and their closed-form structure makes them difficult to apply to large supply network problems.

Due to their optimization-based formulation, model predictive control policies can perform as well as IMC but with the added capability to incorporate constraints. The MPC approach can be readily scaled to multi-node supply chain systems, as shown in Wang et al. (2007) and Schwartz et al. (2006). Comparison of MPC to the smoothing replenishment rule (SRR), an advanced order-up-to policy based on control principles developed by Dejonckheere et al. (2003), demonstrates increased functionality, superior performance and improved flexibility. MPC policies can be tuned to minimize inventory deviation from setpoint, changes in factory starts, or a weighted combination. The paper presented a comparison study which showed that MPC was able to outperform the SRR with respect to both metrics.

Optimal performance of the MPC decision policy hinges on the quality of its exogenous signals, particularly demand forecasts. The detrimental effects of forecast error can be mitigated by applying a forecasting procedure contextualized for the decision policy. Approaches for *control-relevant* demand modeling and forecasting, as shown in Schwartz and Rivera (2009) and Schwartz

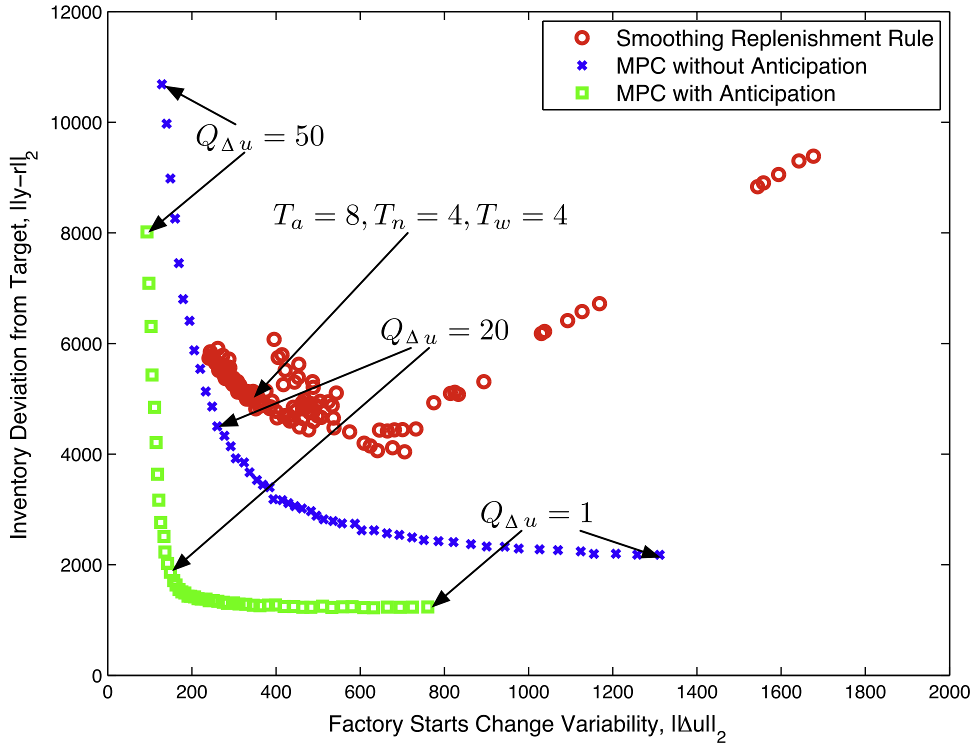

Fig. 14. Comparison of the MPC and SRR policies over a wide range of tuning parameters. Each data point represents an average of 100 stochastic simulations similar to those shown in Fig. 13. Factory thrash is plotted on the *x*-axis; inventory deviation on the *y*-axis. Since improved performance is represented by smaller values the simulation results in the lower-left corner indicate a superior inventory management strategy.

et al. (2009), constitute some of our current ongoing efforts on supply chain management problems. Future research is needed to address the complicated, nonlinear relationship between utilization, throughput, and cycle time present in real manufacturing systems; additionally, the internal model control-based tactical decision policy must be extended to function in multi-echelon supply chains. Approaches to these issues are explored in the Ph.D. dissertation by the first author (Schwartz, 2008).

Acknowledgments

Support from the National Science Foundation (CMMI-0432439) and the Intel Research Council is gratefully acknowledged.

References

Aggelogiannaki, E., Doganis, P., Sarimveis, H., 2008. An adaptive model predictive control configuration for production-inventory systems. International Journal of Production Economics 114, 165–178.
Arrow, K.J., Harris, T., Marschak, J., 1951. Optimal inventory policy. Econometrica 19 (3), 250–272.
Braun, M.W., Rivera, D.E., Flores, M.E., Carlyle, W.M., Kempf, K.G., 2003. A model predictive control framework for robust management of multi-product, multi-echelon demand networks. Annual Reviews in Control 27, 229–245.
Buffa, E.S., Miller, J.G., 1979. Production-Inventory Systems: Planning and Control, third ed. Richard D. Irwin, Inc, Homewood, IL.
Camacho, E.F., Bordons, C., 1999. Model Predictive Control. Springer, London.
Dejonckheere, J., Disney, S.M., Lambrecht, M.R., Towill, D.R., 2002. Transfer function analysis of forecasting induced bullwhip in supply chains. International Journal of Production Economics 78, 133–144.
Dejonckheere, J., Disney, S.M., Lambrecht, M.R., Towill, D.R., 2003. Measuring and avoiding the bullwhip effect: a control theoretic approach. European Journal of Operational Research 147, 567–590.
García, C.E., Prett, D.M., Morari, M., 1989. Model predictive control: theory and practice-a survey. Automatica 25 (3), 335–348.
Glasserman, P., Tayur, S., 1995. Sensitivity analysis for base-stock levels in multiechelon production-inventory systems. Management Science 41, 263–281.
Grubbström, R.W., Wikner, J., 1996. Inventory trigger control policies developed in terms of control theory. International Journal of Production Economics 45, 397–406.
Holt, C.C., Modigliani, F., Muth, J., Simon, H.A., 1960. Planning Production, Inventories, and Work Force. Prentice-Hall, Englewood Cliffs, New Jersey.
Kapuscinski, R., Tayur, S., 1998. A capacitated production-inventory model with periodic demand. Operations Research 46, 899–911.
Kempf, K.G., 2004. Control-oriented approaches to supply chain management in semiconductor manufacturing. In: Proceedings of the 2004 American Control Conference, Boston, MA, pp. 4563–4576.
Lewin, D.R., Scali, C., 1988. Feedforward control in the presence of uncertainty. Industrial Engineering Chemistry Research 27, 2323–2331.
Morari, M., Zafiriou, E., 1988. Robust Process Control. Prentice-Hall, Englewood Cliffs, New Jersey.
Ogunnaike, B.A., Ray, W.H., 1994. Process Dynamics, Modeling, and Control. Oxford University Press, New York.
Ortega, M., Lin, L., 2004. Control theory applications to the production-inventory problem: a review. International Journal of Production Research 42 (11), 2303–2322.
Perea-López, E., Ydstie, B.E., Grossman, I.E., 2003. A model predictive control strategy for supply chain optimization. Computers and Chemical Engineering 27, 1201–1218.
Qin, S.J., Badgwell, T.A., 2003. A survey of industrial model predictive control technology. Control Engineering Practice 11, 733–764.
Riddalls, C.E., Bennett, S., 2002. Production-inventory system controller design and supply chain dynamics. International Journal of Systems Science 33 (3), 181–195.
Rivera, D.E., Flores, M.E., 2004. Internal model control. Vol. Encyclopedia of Life Support Systems of Control Systems, Robotics, and Automation. Eolss Publishers, Oxford, UK 〈http://www.eolss.net〉.
Rivera, D.E., Morari, M., Skogestad, S., 1986. Internal model control 4: PID controller design. Industrial & Engineering Chemistry Process Design and Development 25, 252–265.
Sarimveis, H., Patrinos, P., Tarantilis, C.D., Kiranoudis, C.T., 2008. Dynamic modeling and control of supply chain systems: a review. Computers & Operations Research 35, 3530–3561.
Schwartz, J., 2008. Next generation supply chain management: control, optimization, and system identification. Ph.D. Thesis, Arizona State University.
Schwartz, J., Arahal, M., Rivera, D., Smith, K., 2009. Control-relevant demand forecasting for tactical decision-making in semiconductor manufacturing

supply chain management. IEEE Transactions on Semiconductor Manufacturing 22 (1), 154–163.
Schwartz, J.D., Rivera, D.E., 2006. Control-relevant demand modeling for supply chain management. In: Proceedings of the 14th Annual IFAC Symposium on System Identification, Newcastle, Australia, pp. 267–272.
Schwartz, J.D., Rivera, D.E., 2009. Control-relevant estimation of demand models for closed-loop control of a production-inventory control system. In: Proceedings of the Combined 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference, Shanghai, PRC, pp. 470–475.
Schwartz, J.D., Wang, W., Rivera, D.E., 2006. Simulation-based optimization of model predictive control policies for inventory management in supply chains. Automatica 42 (8), 1311–1320.
Seborg, D.E., Edgar, T.F., Mellichamp, D.A., 2004. Process Dynamics and Control, second ed. Wiley, New York, NY.
Seferlis, P., Giannelos, N.F., 2004. A two-layered optimisation-based control strategy for multi-echelon supply chain networks. Computers and Chemical Engineering 28, 799–809.
Seierstad, A., Sydsæter, K., 1987. Optimal Control Theory with Economic Applications. Elsevier Science, New York, NY.
Sethi, S.P., Thompson, G.L., 2000. Optimal Control Theory: Application to Management Science and Economics, 2nd ed. Kluwer Academic Publishers, Boston, Dordrecht, London.
Thompson, G.L., Sethi, S.P., 1980. Turnpike horizons for production planning. Management Science 26 (3), 229–241.
Towill, D.R., Zhou, L., Disney, S.M., 2007. Reducing the bullwhip effect: looking through the appropriate lens. International Journal of Production Economics 108, 444–453.
Tzafestas, S., Kapsiotis, G., Kyriannakis, E., 1997. Model-based predictive control for generalized production planning problems. Computers in Industry 34, 201–210.
Wagner, H.M., Whitin, T.M., 1958. Dynamic version of the economic lot size model. Management Science 5, 89–96.
Wang, W., 2006. Robust model predictive control strategies for supply chain management in semiconductor manufacturing. Ph.D. Thesis, Arizona State University.
Wang, W., Rivera, D.E., 2008. Model Predictive Control for tactical decision-making in semiconductor manufacturing supply chain management. IEEE Transactions on Control Systems Technology 16 (5), 841–855.
Wang, W., Rivera, D.E., Kempf, K.G., 2007. Model Predictive Control strategies for supply chain management in semiconductor manufacturing. International Journal of Production Economics 107 (1), 56–77.
Warburton, R.D.H., Disney, S.M., 2007. Order and inventory variance amplification: the equivalence of discrete and continuous time analyses. International Journal of Production Economics 110, 128–137.
Whitin, T.M., 1970. The Theory of Inventory Management. Greenwood Press, Westport, CT.
Wilson, R.H., 1934. A scientific routine for stock control. Harvard Business Review 13, 116–128.