

Simulation-based optimization of process control policies for inventory management in supply chains[☆]

Jay D. Schwartz, Wenlin Wang, Daniel E. Rivera*

Control Systems Engineering Laboratory, Department of Chemical and Materials Engineering, Arizona State University, Tempe, AZ 85287-6006, USA

Received 1 April 2005; received in revised form 17 March 2006; accepted 24 March 2006

Available online 9 June 2006

Abstract

A simulation-based optimization framework involving simultaneous perturbation stochastic approximation (SPSA) is presented as a means for optimally specifying parameters of internal model control (IMC) and model predictive control (MPC)-based decision policies for inventory management in supply chains under conditions involving supply and demand uncertainty. The effective use of the SPSA technique serves to enhance the performance and functionality of this class of decision algorithms and is illustrated with case studies involving the simultaneous optimization of controller tuning parameters and safety stock levels for supply chain networks inspired from semiconductor manufacturing. The results of the case studies demonstrate that safety stock levels can be significantly reduced and financial benefits achieved while maintaining satisfactory operating performance in the supply chain.

© 2006 Elsevier Ltd. All rights reserved.

Keywords: Supply chain management; Simulation-based optimization; Inventory control; Internal model control; Model predictive control

1. Introduction

Improved operation of supply chains for manufactured goods is worth billions of dollars to the national economy (Simchi-Levi, Kaminsky, & Simchi-Levi, 2004); effective inventory management plays an important role in this regard. The use of optimization techniques in the management of supply/demand networks began with the development of the classical economic order quantity approach (Wilson, 1934). Later developments include approaches for determining optimal base stock levels in “order-up-to” policies (Glasserman & Tayur, 1995; Kapuscinski & Tayur, 1998) and the application of optimal control theory (Blanchini, Miani, & Rinaldi, 2004; Sethi & Thompson, 2000).

Model predictive control (MPC) (García, Prett, & Morari, 1989) strategies relying on deterministic linear models have recently been proposed for tactical inventory management of

multi-echelon production–inventory systems as seen in supply chains (Braun, Rivera, Flores, Carlyle, & Kempf, 2003; Perea-López, Ydstie, & Grossman, 2003; Seferlis & Giannelos, 2004; Tzafestas, Kapsiotis, & Kyriannakis, 1997). However, the inventory management problems typically found in practice correspond to uncertain, stochastic systems. Consider a representative real-life problem originating from semiconductor manufacturing as shown in Fig. 1. Here finished products (computer chips) are the result of processing silicon wafers through a fabrication/test node, assembling die and packages through an assembly/test node, and finishing the product before shipping to the customer to satisfy orders (Kempf, 2004). This supply chain is characterized by long throughput times, variability in both throughput time and yields, and significant uncertainty in demand. As control-oriented frameworks, internal model control (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.

The primary objective of this work is to present a simulation-based framework for optimally tuning these policies in a stochastic, uncertain environment using the concept of

[☆] This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor Qing Zhang under the direction of Editor Suresh Sethi.

* Corresponding author. Tel.: +1 480 965 9476; fax: +1 480 965 0037.

E-mail addresses: jay.schwartz2@asu.edu (J.D. Schwartz), wenlin.wang@asu.edu (W. Wang), daniel.rivera@asu.edu (D.E. Rivera).

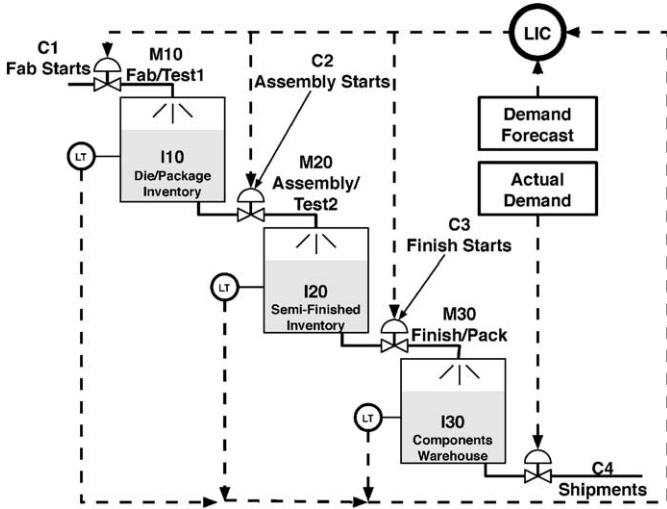


Fig. 1. Fluid representation for a representative three-echelon supply chain based on semiconductor manufacturing (Kempf, 2004).

simultaneous perturbation stochastic approximation (SPSA) (Spall, 2003). SPSA incorporates a simultaneous perturbation optimization method and differs from infinitesimal perturbation analysis (Glasserman & Tayur, 1995) and the Robbins–Monro stochastic approximation algorithm (Robbins & Monro, 1951) in that it avoids an explicit calculation or measurement of the gradient (Spall, 1998, 2003).

Two important scenarios are presented that illustrate the benefits of the SPSA approach in enhancing the usefulness of the policies under conditions of uncertainty. In the first scenario, an IMC decision policy for a single product, single echelon production–inventory system is evaluated. The SPSA technique is applied to determine financially optimal controller tuning parameters under conditions involving varying magnitudes of forecast error. The second problem scenario involves the simultaneous selection of safety stock targets and MPC move suppression parameters for the representative semiconductor manufacturing problem described in Fig. 1 in circumstances involving stochastic yield, variable throughput times, and uncertain, autocorrelated demand.

The paper is organized as follows. Sections 2 and 3 are concerned with presenting the IMC- and MPC-based tactical decision policies that are optimal with respect to linear time-invariant models derived using fluid analogies. Section 4 describes the SPSA optimization method that will seek optimal tuning and targets of the IMC and MPC policies when placed in a stochastic environment. Section 5 presents the results of applying SPSA for the previously described scenarios, with the results yielding some fundamental insights into the proper selection of inventory targets and tuning of the decision policies. A summary of the work and resulting conclusions is presented in Section 6.

2. Multi-degree-of-freedom combined feedback–feedforward internal model control

Fluid analogies represent meaningful descriptions of supply chains associated with high-volume manufacturing problems

at sufficiently long time scales. 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 (retailers, distributors, etc.). The warehouse serves as a buffer in the presence of stochastic, uncertain customer demand and factory output. 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. 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 equation

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

Based on (1) it is possible to derive feedback-only decision policies that manipulate factory starts to maintain inventory level at a designated setpoint. However, if knowledge of future customer demand is available, it is advantageous to use 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

$$d(t) = d_F(t - \theta_F) + d_U(t). \quad (2)$$

The overall dynamical system is then defined by the equations:

$$\begin{aligned} y(s) &= p(s)u(s) - p_{d1}(s)p_{d2}(s)d_F(s) - p_{d2}(s)d_U(s) \\ &= \frac{Ke^{-\theta s}}{s}u(s) - \frac{e^{-\theta_F s}}{s}d_F(s) - \frac{1}{s}d_U(s). \end{aligned} \quad (3) \quad (4)$$

The model per Eq. (4) is the basis for the process control-based tactical decision policies considered in this paper. We first evaluate a multi-degree-of-freedom combined feedback–feedforward IMC structure (Morari & 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. 2 shows the structure schematically. The aforementioned controllers correspond to q_r for setpoint tracking, q_F for measured disturbance rejection, and q_d for unmeasured disturbance rejection.

The IMC design procedure for these controllers is comprised of 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,

$$\begin{aligned} \min_{\tilde{q}_r} \|(1 - \tilde{p}\tilde{q}_r) r\|_2, \\ \min_{\tilde{q}_d} \|(1 - \tilde{p}\tilde{q}_d)p_{d2} d_U\|_2, \\ \min_{\tilde{q}_F} \|(\tilde{p}_d - \tilde{p}\tilde{q}_F)p_{d1} p_{d2} d_F\|_2, \end{aligned}$$

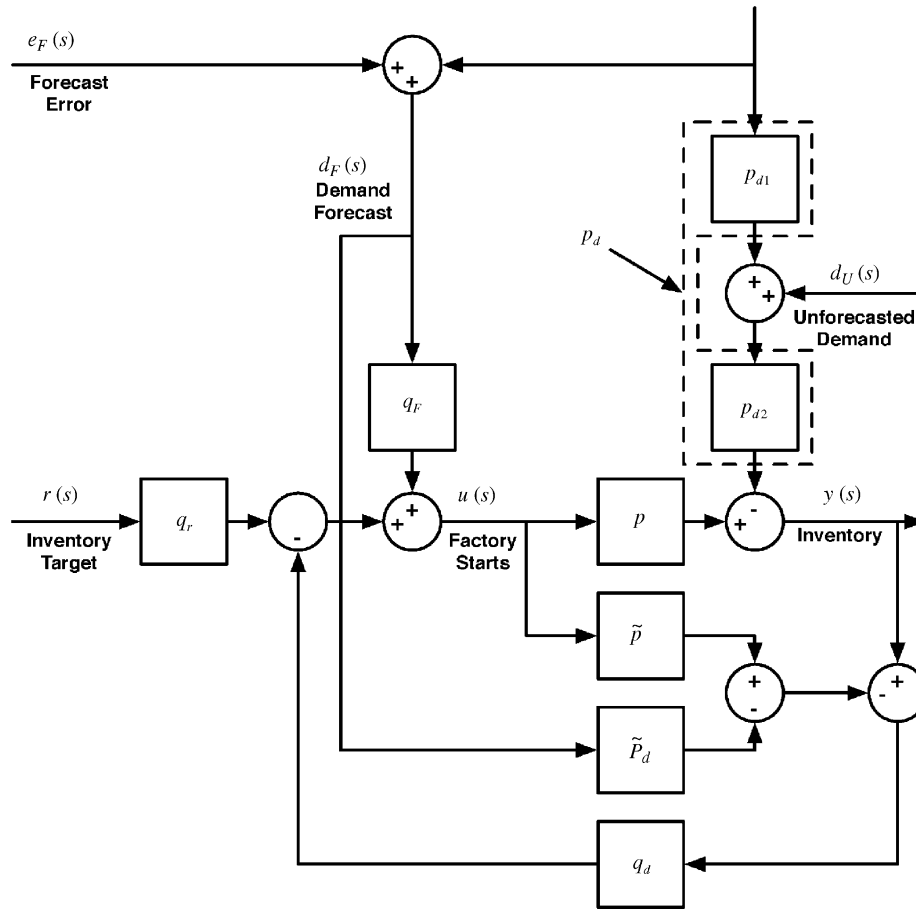


Fig. 2. Three-degree-of-freedom combined feedback-feedforward IMC structure.

subject to the requirements that $\tilde{q}_r(s)$, $\tilde{q}_d(s)$ and $\tilde{q}_F(s)$ be stable and causal.

- (2) *Design for robust stability and performance:* $\tilde{q}_r(s)$, $\tilde{q}_d(s)$ and $\tilde{q}_F(s)$ are augmented with 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 are shown as follows:

Setpoint tracking: The setpoint tracking mode of this control system is designed for H_2 -optimality with respect to step inputs and augmented with a lowpass filter. This controller guarantees no offset for Type-1 setpoint changes in the control system.

$$q_r(s) = \frac{s}{(\lambda_r s + 1)^{n_r}} \quad (5)$$

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 factorization for ramp inputs, with a Type-2 filter

guaranteeing no offset for both asymptotically step and ramp disturbances.

$$q_d(s) = s(\theta s + 1) \frac{(n_d \lambda_d s + 1)}{(\lambda_d s + 1)^{n_d}} \quad (6)$$

Measured disturbance rejection: The measured disturbance rejection mode relies on a θ_F -day ahead forecast signal to manipulate factory starts. The IMC controller form is defined as follows (Lewin & Scali, 1988)

$$q_F(s) = q'_F(s) f_F(s), \quad (7)$$

where $q'_F(s)$ is defined as

$$q'_F(s) = e^{-(\theta_F - \theta)s} \quad (8)$$

if the forecast horizon is longer than the factory throughput time ($\theta_F \geq \theta$). If the forecast horizon is shorter ($\theta_F < \theta$) then $q'_F(s)$ is defined as

$$q'_F(s) = (\theta - \theta_F)s + 1. \quad (9)$$

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

$$f_F(s) = \frac{(n_F \lambda_F s + 1)}{(\lambda_F s + 1)^{n_F}} \quad (10)$$




Fig. 3. IMC system response to a unit pulse in forecast error ($\theta = 2$, $\theta_F = 5$, $\lambda_d = 2$, $\lambda_F = 2$, $K = 1$, $n_d = n_F = 3$), note that the bottom plot shows the θ_F -day ahead forecast error.

Each controller is required to be stable and proper, thus imposing the constraint that all values of the user-adjustable parameters (λ_r , λ_d , and λ_F) be positive and that the filter order is chosen to ensure transfer function properness ($n_r \geq 1$, $n_d \geq 3$, $n_F \geq 3$). Fig. 3 shows some representative results for the response of the IMC control system to a forecast error pulse. The controller anticipates the increased future demand and increases starts accordingly. When no demand change is realized, starts are decreased to return the inventory level to the setpoint.

3. Model predictive control as a tactical decision policy

The application of MPC to the inventory management problems considered in this paper follows along the lines of the conceptual framework presented by Braun et al. (2003); this is summarized schematically in Fig. 4. Factory starts are optimized over a move horizon to minimize deviations from inventory targets given an anticipated demand signal. The starts level corresponding to the first entry in the move horizon is implemented and the process is repeated. A meaningful objective function formulation is as follows:

$$\min_{\Delta u(k|k) \dots \Delta u(k+M-1|k)} J, \tag{11}$$

where $\Delta u(k) \dots \Delta u(k + M - 1)$ represents the computed sequence of starts changes and the individual terms of J

correspond to:

$$\begin{aligned}
 & \text{Keep Inventories at Inventory Planning Setpoints} \\
 & \overbrace{\sum_{\ell=1}^P Q_e(\ell) (\hat{y}(k + \ell|k) - r(k + \ell))^2} \\
 & \text{Penalize Changes in Starts} \\
 & + \overbrace{\sum_{\ell=1}^M Q_{\Delta u}(\ell) (\Delta u(k + \ell - 1|k))^2} \tag{12}
 \end{aligned}$$

subject to constraints on inventory capacity ($0 \leq y(k) \leq y_{\max}$), factory inflow capacity ($0 \leq u(k) \leq u_{\max}$), and changes in the quantity of factory starts ($\Delta u_{\min} \leq \Delta u_k \leq \Delta u_{\max}$). Eq. (12) is a multi-objective expression that addresses the main operational objectives in the supply chain. The first term is a setpoint tracking term intended to maintain inventory levels at user-specified targets over time. The second term is a move suppression term that penalizes changes in the factory starts. The emphasis given to each one of the sub-objectives in (12) (or to specific system variables within these objective terms) is achieved through the choice of weights ($Q_e(\ell)$ and $Q_{\Delta u}(\ell)$). These can potentially vary over the move and prediction horizons (M and P , respectively).

Controlled variables y for the problem in Fig. 1 consist of the three inventory levels (I_{10} , I_{20} , and I_{30}). The starts rates for the Fab/Test1, Assembly/Test2, and Finish nodes (C_1 , C_2 and C_3) represent manipulated variables. The demand signal (which

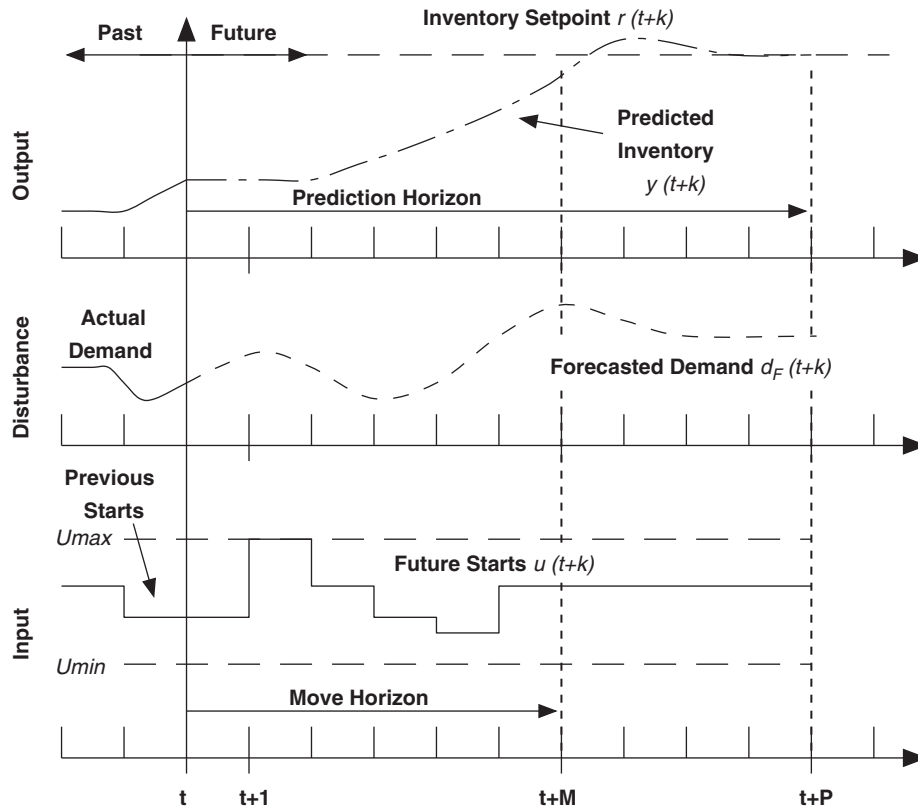


Fig. 4. Receding horizon representation of MPC.

dictates the shipment flow in C_4) is treated as a disturbance and consists of two components: (1) actual demand (which is only fully known after the fact) and (2) forecasted demand, which is provided to the algorithm as an exogeneous signal. For the problem in Fig. 1 the mass conservation relationship for die-package inventory (I_{10}) can be written as

$$I_{10}(k + 1) = I_{10}(k) + Y_1 C_1(k - \theta_1) - C_2(k), \quad (13)$$

θ_1 and Y_1 represent the nominal throughput time and yield for the Fab/Test1 node, respectively, while C_1 and C_2 represent the daily (or per-shift) starts that constitute inflow and outflow streams for I_{10} and M_{10} . Similar relationships to (13) can be written for the semi-finished goods (I_{20}) and components warehouse (I_{30}) inventories. These material balances can in general be organized into a discrete-time state-space model representation amenable to MPC implementation and analysis.

4. Simulation-based optimization using SPSA

Simulation-based optimization algorithms are generally applied to problems where a closed-form relationship between the parameters being optimized and the objective function is unknown or difficult to obtain. This may be due to the presence of noise in the objective function evaluation, or the relationship between the parameters and the function may be significantly complex. The lack of explicit gradient information justifies an interest in optimization algorithms that rely solely on measurements of the objective function.

Sections 2 and 3 showed the development of controllers that are based on nominal linear models, but will be implemented in uncertain, stochastic settings. To achieve optimality in a stochastic setting, it is necessary to add a second optimization layer. A simulation-based optimization scheme is particularly useful in this regard. In this paper we examine SPSA, a promising approach that has received considerable attention over the last decade (Spall, 1998, 2003). The method has been used in a wide variety of settings such as statistical parameter estimation, adaptive control, and many other applications (Spall, 2003).

The SPSA technique represents a significant improvement over traditional finite-difference stochastic approximation (FDSA) methods. The basis of the method is an efficient and intuitive “simultaneous perturbation” estimate of the gradient. Only two measurements of the objective function are required at each iteration, regardless of the number of parameters p . SPSA realizes the same level of accuracy as comparable FDSA methods for a given number of iterations despite the fact that only two measurements are made to form an estimate, as opposed to $2p$ measurements. Therefore, SPSA requires p times fewer evaluations of the objective function to achieve an equivalent result (Spall, 2003).

The underlying premise of SPSA is the minimization of an objective function, J . The objective function J takes a real-valued vector of search parameters \vec{x} of dimension p and returns a scalar. The process begins with an initial guess of the input vector \vec{x} and iterates using the simultaneous perturbation estimate of the gradient $g(\vec{x}) = \delta J / \delta \vec{x}$. Note that this

formulation is similar to the FDSA algorithm discussed previously, but differs in the nature of the gradient estimate. The SPSA algorithm consists of the following steps:

- (1) *Initialize the input vector and gain sequences:* An initial guess of the optimal input vector is made (\vec{x}_0). At this stage, one must also select the coefficients of the gain sequences $a_k = f(a, \alpha)$ and $c_k = f(c, \gamma)$. These sequences govern the step size at each iteration and the magnitude of the perturbation, respectively. Guidelines for generating the gain sequence coefficients a and c are available in the literature (Spall, 1998, 2003). Values for α and γ used in the case studies of the next section (0.602 and 0.101, respectively) were obtained directly from Section III in Spall, 1998.
- (2) *Generate the perturbation vector:* A random perturbation vector (Δ_k) is generated. Each element of the vector is independently generated using a Bernoulli ± 1 distribution with a probability of $\frac{1}{2}$ for each possible outcome.
- (3) *Evaluate the objective function:* Two measurements of the objective function are obtained: $J(\vec{x}_k + c_k \Delta_k)$ and $J(\vec{x}_k - c_k \Delta_k)$.
- (4) *Approximate the gradient:* The simultaneous perturbation approximation of the gradient, $\hat{J}(\vec{x}_k)$, is determined using

$$\hat{g}(\vec{x}_k) = \frac{J(\vec{x}_k + c_k \Delta_k) - J(\vec{x}_k - c_k \Delta_k)}{2c_k \Delta_k}. \quad (14)$$

- (5) *Update the estimate:* The standard stochastic approximation form (15) is used to update \vec{x}_k to \vec{x}_{k+1} .

$$\vec{x}_{k+1} = \vec{x}_k - a_k \hat{g}(\vec{x}_k). \quad (15)$$

5. Case studies

In both case studies described in this section we consider a financial objective function as the basis for optimizing the decision policy. Profit is determined according to the objective function

$$\max_{\vec{x}} \text{Profit} = \text{Revenue} - \text{Production}_{\text{cost}} - \text{Inventory}_{\text{cost}} - \text{Backorder}_{\text{cost}}, \quad (16)$$

where (for the MPC case study)

$$\text{Revenue} = \sum_{k=1}^{T_{\text{final}}} \gamma_R C_4(k), \quad (17)$$

$$\text{Production}_{\text{cost}} = \sum_{k=1}^{T_{\text{final}}} \sum_{j=1}^{N_{\text{nodes}}} \gamma_{C_j} C_j(k), \quad (18)$$

$$\text{Inventory}_{\text{cost}} = \sum_{k=1}^{T_{\text{final}}} \sum_{j=1}^{N_{\text{nodes}}} \gamma_I I_{10j}(k), \quad (19)$$

$$\text{Backorder}_{\text{cost}} = \sum_{k=1}^{T_{\text{final}}} \gamma_B \text{Backorders}(k). \quad (20)$$

An equivalent objective function is formulated for the continuous time production–inventory system by sampling daily. The search vector \vec{x} consists of important system parameters that could include IMC controller tuning parameters (λ_r , λ_d , and λ_F), inventory targets (r_i), objective function weights ($Q_e(\ell)$, $Q_{\Delta u}(\ell)$), nominal model parameters, prediction horizons, and so forth. The objective function shown in (16) comprehensively accounts for the production cost, inventory holding cost, backorder penalty, and revenue generation via the parameters γ_C , γ_I , γ_B , and γ_R , respectively.

5.1. Case study 1: tuning a single echelon supply chain

Consider a production–inventory system with deterministic supply ($\theta = 2$, $\theta_F = 5$, $K = 1$) and uncertain, autoregressive demand (generated by filtering a white noise sequence with a first-order lowpass filter with bandwidth of 1 rad/day). Forecast error corresponds to a white noise signal. An IMC decision policy as described in Section 2 is developed based on these nominal values. Weights for the financial objective function are $\gamma_R = 5$, $\gamma_{C_1} = 2$, $\gamma_I = 0.1$, and $\gamma_B = 5$. Fig. 5 shows optimal values of the adjustable-parameters λ_d and λ_F for increasing values of forecast error variance. When forecast error is low, the feedforward controller can be tuned aggressively, which is reflected in low values of λ_F . As the amount of forecast error increases, the SPSA search indicates that financial optimality is attained by implementing a more passive feedforward controller in conjunction with an increasingly aggressive feedback controller. For high levels of forecast error it is advantageous to act cautiously to forecast information, but be more aggressive with feedback action.

5.2. Case study 2: optimizing inventory targets and controller tuning in a three echelon supply chain

For the second case study, the SPSA algorithm is applied to the larger network topology ($N_{\text{nodes}} = 3$) shown in Fig. 1 subject to a centralized MPC decision policy. The throughput time in M_{10} varies according to a triangular distribution, with 80% of the output produced after 35 days and the remaining 20% evenly split between days 34 and 36. Throughput times vary similarly for M_{20} and M_{30} , with ranges between 5–7 and 1–3 days, respectively. Yield rates vary uniformly for each manufacturing node ($95\% \pm 2\%$ for M_{10} , $98\% \pm 2\%$ for M_{20} , and $99\% \pm 1\%$ for M_{30}). Stochasticity and uncertainty in demand are considered, as shown in Fig. 6, which shows a demand realization from an autoregressive process and its corresponding forecast. The MPC prediction horizon P is 70 days and the move calculation horizon M is 60 days. In all cases, the emphasis is on meeting the inventory targets equally ($Q_e(\ell) = 1$) for all channels and elements of the horizon. Weights for the financial objective function are as follows: $\gamma_R = 40$, $\gamma_{C_1} = 10$, $\gamma_{C_2} = 8$, $\gamma_{C_3} = 2$, $\gamma_I = 0.1$, and $\gamma_B = 5$. Fig. 7 shows the optimization path for the second case study. Initially, the optimizer drives the inventory targets towards zero. As baseline inventory




Fig. 5. Financially optimal IMC controller tuning parameters (λ_d and λ_F) vs. forecast error variance.




Fig. 6. Actual and forecasted customer demand for the MPC case study.

levels are decreased, the risk of backorders increases. Eventually, the optimization algorithm converges to an optimum where the inventory targets are approximately 400, 0, and 1000 units

for I_{10} , I_{20} , and I_{30} , respectively. Given supply and demand uncertainty, it makes physical sense to keep a buffer in the final inventory stage (I_{30} , the inventory closest to the demand), but




Fig. 7. SPSA optimization algorithm path for a six-dimensional search space involving the selection of inventory targets and move suppression values. Top: profit, Middle: inventory targets, Bottom: move suppression values.

seek to minimize the amount of excess intermediate products stored in earlier stages. Note that the optimal target for I_{10} is greater than that of I_{20} , this agrees with intuition as the factory M_{10} has the longest throughput time and most stochastic behavior of all the production nodes. As demand variability and stochasticity increases, it will become necessary to keep larger inventories at all levels of the supply/demand network. For the cost values used in this study and the characteristics of the demand signal, profitability of the supply chain seems somewhat less sensitive to changes in the move suppression values than to changes in setpoint targets. Inventory and backorder costs are significantly greater than costs incurred by detuning the MPC decision policy.

Fig. 8 shows the MPC simulation where controller tuning and inventory targets are determined from the SPSA optimization algorithm (the final iteration shown in Fig. 7). Safety stock levels are reduced to a level where inventory targets are as low as possible without incurring backorders, and increased move suppression denotes less “thrash” in factory starts. These conditions minimize inventory holding costs and increase profits.

6. Summary and conclusions

IMC and MPC algorithms have been demonstrated to be capable of managing inventory in uncertain production–inventory and multi-echelon supply/demand networks, respectively. The use of SPSA allows for determining controller

tunings and operating targets that lead to optimal results from either an operational or financial standpoint. The results of the optimization on a single node example show that it is advantageous to act cautiously to forecasted information and gradually become more aggressive (with respect to feedforward action) as more accurate demand forecasts become available. For the three echelon problem, the use of the simulation-based optimization method led to insights concerning the proper parameterization and tuning of the tactical MPC decision policy. The amount of safety stock necessary for optimal profitability is a function of the accuracy and magnitude of the demand forecast. SPSA provides a way of systematically determining the financially optimal inventory targets and the move suppression values present in the MPC objective function simultaneously. For the semiconductor manufacturing problem case study, it was found that the optimization problem was more sensitive to changes in inventory targets, and less sensitive to changes in move suppression. This allows for flexibility when tuning the decision policy, as robustness considerations do not have to be cast aside in favor of increased profitability.

Acknowledgements

The authors would like to acknowledge support from the Intel Research Council and the National Science Foundation (DMI-0432439). Dr. Karl Kempf provided valuable insights that facilitated the research.

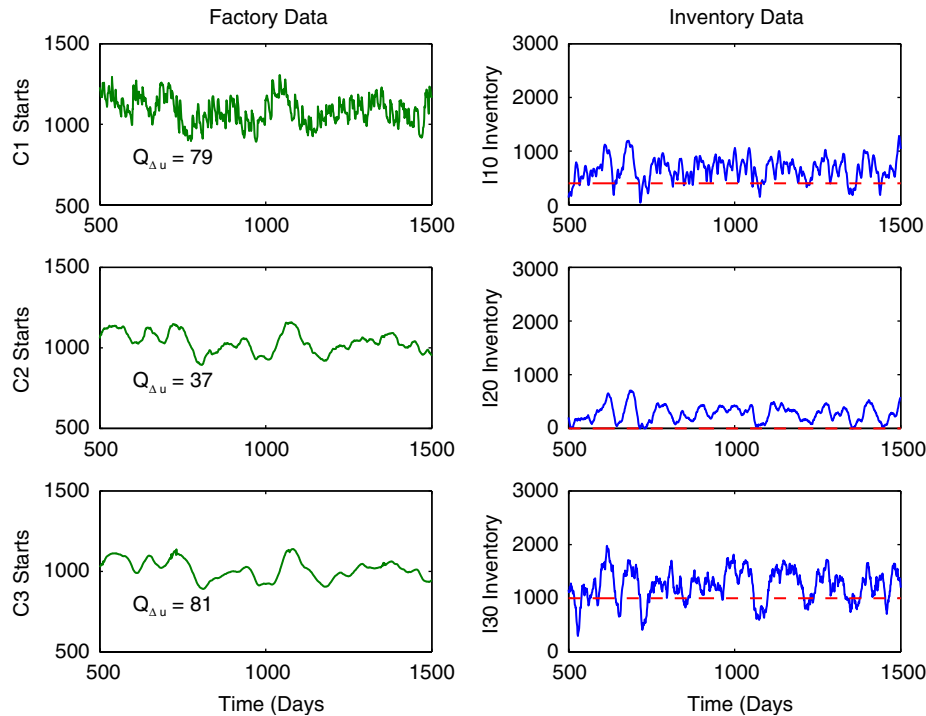


Fig. 8. Time series for the three echelon supply chain shown in Fig. 1 with MPC move suppression and setpoint targets chosen via the SPSA optimization strategy.

References

- Blanchini, F., Miani, S., & Rinaldi, F. (2004). Guaranteed cost control for multi-inventory systems with uncertain demand. *Automatica*, 40(2), 213–223.
- 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.
- 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.
- Kapusinski, 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*. (pp. 4563–4576), Boston, MA.
- 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*. Englewood Cliffs, NJ: Prentice-Hall.
- 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.
- Robbins, H., & Monro, S. (1951). A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3), 400–407.
- 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.
- Sethi, S. P., & Thompson, G. L. (2000). *Optimal control theory: Application to management science and economics*. 2nd ed., Boston/Dordrecht/London: Kluwer Academic Publishers.
- Simchi-Levi, D., Kaminsky, P., & Simchi-Levi, E. (2004). *Managing the supply chain*. New York, NY: McGraw-Hill.
- Spall, J. C. (1998). Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on Aerospace and Electronic Systems*, 34(3), 817–823.
- Spall, J. C. (2003). *Introduction to stochastic search and optimization estimation, simulation, and control*. Hoboken NJ: Wiley.
- Tzafestas, S., Kapsiotis, G., & Kyriannakis, E. (1997). Model-based predictive control for generalized production planning problems. *Computers in Industry*, 34, 201–210.
- Wilson, R. H. (1934). A scientific routine for stock control. *Harvard Business Review*, 13, 116–128.



Jay D. Schwartz received the B.S.E. degree in Chemical Engineering in 2004 from Arizona State University, where he is currently pursuing the Ph.D. degree. His research interests include predictive control, numerical optimization, and system identification as related to supply chain management.



Wenlin Wang received his B.E. in Electrical Engineering from Hefei University of Technology, China in 1998, an M.S. in Control Science and Engineering from Zhejiang University in 2001 and an M.S.E in Chemical Engineering from Arizona State University. He is currently a Ph.D. candidate in Chemical Engineering at Arizona State University. His research interests include model predictive control, advanced process control and supply chain management in semiconductor manufacturing.



Daniel E. Rivera is Associate Professor in the Department of Chemical and Materials Engineering at Arizona State University, and Program Director for the ASU Control Systems Engineering Laboratory. Prior to joining ASU he was a member of the Control Systems Section of Shell Development Company in Houston, Texas. He received his Ph.D. in chemical engineering from the California Institute of Technology and holds B.S. and M.S. degrees from the University of Rochester and the University of Wisconsin-Madison, respectively. He has been

a visiting researcher with the Division of Automatic Control at Linköping University, Sweden, Honeywell Technology Center in Phoenix, AZ, and the University of St. Cyril and Methodius in Skopje, Macedonia. His research

interests include the topics of system identification, robust process control, and the application of control ideas to “non-traditional” problems such as inventory management in enterprise systems and adaptive interventions in behavioral health.

Dr. Rivera is a member of the IFAC Technical Committee in Modeling, Identification and Signal Processing and is also an Associate Editor for the IEEE Transactions in Control Systems Technology and the IEEE Control Systems Magazine. He is the recipient of the 1994–1995 Outstanding Undergraduate Educator Award by the ASU student chapter of AIChE and the 1997–1998 Teaching Excellence Award from the ASU College of Engineering and Applied Sciences.