COMPUTATION OF OPTIMAL FLOW CONTROL POLICIES OF A MANUFACTURING SYSTEM WITH MULTIPLE PRODUCTION RATES

(preliminary, condensed version)

Gül Gürkan
CentER for Economic Research
Tilburg University
P.O. Box 90153
5000 LE Tilburg, The Netherlands
email: ggurkan@kub.nl

Fikri Karaesmen
Laboratoire d'Informatique de Paris 6
Universite Pierre et Marie Curie
4, place Jussieu
75252 Paris Cedex 05, France
email: Fikri.Karaesmen@lip6.fr

KEYWORDS

Manufacturing flow control, Stochastic optimization, Sensitivity analysis, Generalized semi-Markov process, Hedging point control policies.

ABSTRACT

We address the problem of finding optimal hedging points of a production flow controller, in the sense of Bielecki and Kumar (1988). We model the dynamics of the controller by using generalized semi-Markov processes (GSMP); the GSMP framework lets us work with multiple states, and not only with exponential distribution but with fairly general distributions. Utilizing the GSMP representation, the sensitivity estimates of cost functions, consisting of inventory and backlog costs, with respect to hedging points are computed. We then discuss how these sensitivity estimates are used together with sample-path optimization, a recent simulation-based optimization method, for finding optimal hedging points.

1 INTRODUCTION

Ever since the pioneering work of Kimemia and Gershwin (1983), there has been considerable interest in modeling and resolving production control problems as optimal control problems for stochastic flow systems. This *manufacturing flow control* framework assumes that, the production/inventory processes of interest are governed by deterministic flows that correspond to demand and production rates and that the randomness in the system can be captured by changes in production or demand rates according to some stochastic process. For simplicity of exposition, the production process can be thought of being generated by one or more identical machines that can process items at a deterministic rate and the demand is a deterministic flow. The random disturbances to the system can then be viewed as changes in the production capacity of the machine (due typically to a machine breakdown and eventual repair). The flow control problem for this system is to dynamically select the production rate of the machine such that the inventory and backorder costs are minimized.

The particular manufacturing flow control problem that we consider is that of a single machine coping with a deterministic flow of demand. The production capacity of the machine can take multiple values according to some stochastic process and the objective is to control the rate of the machine (within the capacity constraints) to minimize inventory and backorder costs. Our goal in this paper is not to solve the optimal control problem in full generality which would necessitate the determination of the structure of an *optimal policy*, which is a highly challenging problem. Instead, we assume a plausible class of policies which may not be optimal in all cases but is feasible and interesting for implementation.

The essence of our manufacturing flow control problem is the trade-off between inventory and backorder costs. As first discussed by Kimemia and Gershwin (1983), this motivates the idea of *hedging point* control. The system should be controlled such that the controller drives the system to some determined positive inventory level (i.e. the hedging point) whenever possible to hedge against future breakdowns that may cause eventual backorder costs (which are normally higher than inventory holding costs).

For a single machine that can take two maximum production levels (on or off), the optimality of a hedging point policy was theoretically justified by the results of Akella and Kumar (1986) and Bielecki and Kumar (1989) when breakdown and repair intervals are exponentially distributed. For the problem that we consider here, the hedging point control framework suggests multiple hedging points that depend on the production capacity level. There is strong evidence that, this is indeed the optimal policy when production capacity change intervals are exponentially distributed. For non-exponential distributions the optimal policy is likely not to be of hedging point type as the optimal production rate should also depend on the time elapsed since the last production capacity level change. However, even in this case, hedging point policies are practically very interesting since they are very easy to implement. Note that, a hedging point policy is described completely by the specification of a few parameters corresponding to hedging points.

Really relevant literature includes: Brémaud *et al.* (1997), Caramanis and Liberopoulos (1992), Liberopoulos and Caramanis (1994), Haurie *et al.* (1994), Yan *et al.* (1994).

Loosely related literature includes: Hu and Xiang (1993), Sharifnia (1988), ...

Our primary objective is to develop an efficient and robust numerical method that can be used to compute optimal hedging points for a system which is subject to fairly general random disturbances (i.e. not necessarily exponentially distributed). We will use a novel methodology, discrete-event simulation coupled by sample-path optimization.

We first model the dynamics of the system by using a generalized semi-Markov process (GSMP); the GSMP framework lets us work with multiple states and with fairly general distributions. A GSMP can be thought as a mathematical framework which models the evolution of a discrete-event simulation. An excellent description of this framework can be found in Shedler (1993). The basic idea of a GSMP can be briefly explained as follows: There is a set of states and a set of events. The GSMP jumps from one state to another upon the occurrence of an event; at each state there are some active events. At any time, each active event is associated with a clock representing the residual lifetime of that event and a speed at which the clock runs down. If the clock corresponding to event e in state s equals k and the speed at which this clock runs is r, then e is scheduled to occur after k/r units of time. The next event and the time until it occurs are always determined by the smallest clock reading/clock speed ratio. Upon the occurrence of an event, changes may occur in the physical state, clock settings, and clock speeds. By precisely specifying all these entities and how they interact with each other over time, one can model the evolution of an entire system through some recursive expressions. Although, in the past GSMP's were mainly used to model systems with discrete entities, they can also be used to model fluid systems; see Suri and Fu (1994) for example.

Utilizing our GSMP representation, we compute the sensitivity estimates of cost functions, consisting of inventory and backlog costs, with respect to the hedging points. We then use these sensitivity estimates together with *sample-path optimization* for finding optimal hedging points.

Sample-path optimization is a recent simulation-based optimization method to optimize performance functions of complex stochastic systems; see Gürkan *et al.* (1998) for a brief overview of the so-called "sample-path methods" that can be used for providing solutions for difficult stochastic optimization problems and stochastic variational inequalities. The idea is to observe a fixed sample path (by using the method of common random numbers from the simulation literature), solve the resulting deterministic problem using fast and effective methods from nonlinear programming, and then use the resulting solutions to infer information about the solution of the original stochastic problem. In Section 3, we briefly explain the basic ideas and give references in which more details can be found.

2 PROBLEM DESCRIPTION

Consider a manufacturing system that produces a single product and a deterministic demand rate d. The maximum production capacity of the system can take one of K possible values, r_i $(i = 1, 2, \ldots, K)$ according to a stationary process. We do not make any distributional assumptions about the underlying distributions except from requiring them to be non-deterministic and to have support on $(0, \infty)$. (At some point, we may also need to assume that they are independent of each other.)

We denote by $\alpha(t)$ $(\alpha(t) = 1, 2, ..., K)$ the state of the machine at time t and $r_{\alpha(t)}$ the corresponding maximum production capacity in that state. Let $X(t)$ denote the inventory level at time t. Note that, $X(t)$ can take negative values corresponding to backorders in the system. To this end, it is useful to distinguish the positive and negative parts of the inventory process. We define by $X^+(t) = X(t) \cdot I_{\{X(t) \geq 0\}}$ the production surplus of the system and by $X^-(t) = -X(t) \cdot I_{\{X(t)<0\}}$ the backorder level (where I_A is the indicator function corresponding to set A). The instantaneous cost function is then given as: $hX^+(t) + bX^-(t)$ where h and b are respectively the holding and backorder costs (per item per unit time).

We can classify the states of the machine depending on the maximum production rates. A machine state i is called deficient if $r_i < 0$ (if $\alpha(t) = i$, $X(t)$ has to be decreasing regardless of the control selected). Similarly, state i is called non-deficient if $r_i > 0$, in this case $X(t)$ can be increasing, decreasing and constant depending on the actual control. Finally, if $r_i = d$, state i is called a zero machine state and $X(t)$ is either decreasing or constant.

Let $u(t)$ ($r_{\alpha(t)} \geq u(t) \geq 0$) denote the actual production rate at time t. The objective of the manufacturing flow control problem is to select $u(t)$ such that:

$$E_{u,x}\left[\lim_{T\to\infty}\frac{1}{T}\int_0^T hX^+(t) + bX^-(t)dt\right] \qquad (2.1)$$

(where $X(0) = x$) is minimized.

We study the above problem with the restriction that $u(t)$ is in the class of hedging point policies. A hedging point policy drives the system to the hedging point of the current machine state at the fastest rate possible and then keeps it at the hedging point by adjusting the production rate until a machine state change. Accordingly, for each non-deficient or zero machine state i, there is a hedging point z_i. The control $u(t)$ is completely determined by the specification of z_i's for non-deficient or zero machine states:

$$u(t) = \begin{cases} r_i & \text{if } X(t) < z_i \\ d & \text{if } X(t) = z_i \\ 0 & \text{if } X(t) > z_i \end{cases}$$

and for deficient states $u(t) = r_i$ regardless of the inventory position.

The control problem under the restriction to this particular class of policies is to select the optimal values of the hedging points, $(z_1, \ldots, z_K)$, in order to minimize (2.1).

An important question about this system is stability. Let $\pi(i)$ $(i = 1, 2, .., K)$ denote the stationary probability of being in machine state i. Brémaud *et al.* (1997) show that a sufficient condition for stability of the system is:

$$\sum_{i=1}^{K} \pi(i) r_i > d$$

Here on, we will restrict our attention to systems satisfying the above condition.

In Gürkan and Karaesmen (1998), we rigorously develop a GSMP representation to describe the dynamics and to simulate this system. As mentioned earlier, developing the GSMP representation involves first specifying set of states, set of events, clock reading and clock speeds associated with active events at each state, and then describing how events occur and what changes take place in the physical state, clock settings and clock speeds upon the occurrence of events. This eventually leads to some recursive expressions modeling the evolution of the entire system through time. Since the GSMP representation is rather detailed, we do not include it here and refer the reader to Gürkan and Karaesmen (1998).

3 SAMPLE-PATH OPTIMIZATION

In this section we describe the sample-path optimization method in the form proposed in Plambeck *et al.* (1996) and analyzed in Robinson (1996). The method is for minimizing a function f_∞ that is an almost-sure limit of a computable sequence of random functions f_n. In most cases, we have in mind using discrete-event simulation to observe the f_n's. In systems that evolve over time, we simulate operation of the system for n, say, time units and then compute an appropriate performance measure. To observe f_n at different parameter settings, we use the method of common random numbers. Furthermore, in many cases derivatives or directional derivatives of the f_n can be obtained using well-established methods of gradient estimation such as infinitesimal perturbation analysis (IPA); see Ho and Cao (1991) and Glasserman (1991).

Many problems in simulation optimization can be formally modeled by an extended-real-valued stochastic process $\{f_n(x) \mid n = 1, 2, \ldots\}$. The f_n take values that may be real numbers or $\pm\infty$, whereas the parameter x takes values in $\mathbf{R}^k$. Using extended-real-valued random variables is a convenient way of modeling deterministic constraints, since we can always set $f_n(x) = +\infty$ for those x that do not satisfy the constraints. We generally require the functions never to be $-\infty$ and not to be everywhere $+\infty$ (meaning there is at least one feasible point, x). For each $n \geq 1$ and each $x \in \mathbf{R}^k$, $f_n(x)$ is a random variable defined on a common probability space $(\Omega, \mathcal{F}, \mathcal{P})$. In the following we refer to $f_n(x)$ as the sample function and we write $f_n(\omega, x)$ when we want to emphasize the dependence of $f_n(x)$ on the sample point ω.

We assume the existence of a limit function f_∞ such that the f_n almost surely converge pointwise to f_∞ as $n \to \infty$. For the systems we are concerned with, such existence and convergence can often be inferred from regeneration theorems and/or the strong law of large numbers.

As mentioned earlier, we are interested in finding the infimum and, if it exists, a minimizer of f_∞. In general we can only observe f_n for finite n. Therefore we will approximate minimizers of f_∞ using such information about f_n. The method is simple: we fix $\omega \in \Omega$ and compute a minimizer $x_n^*(\omega)$ of $f_n(\omega, .)$ for large n, then take $x_n^*(\omega)$ as an approximate minimizer of $f_\infty(\omega, .)$. Of course, minimizers of $f_\infty(\omega, \cdot)$ will generally depend on the sample point ω. However, in many practical problems in which one would anticipate using this technique f_∞ is a deterministic function, for example a steady-state performance function or an expected value; *i.e.,* it is independent of ω. See Robinson (1996) for a precise statement of convergence results, proofs, and a review of similar ideas that have appeared in the literature; see also Gürkan *et al.* (1998) for a review of applications in which these ideas have been successfully used, and an extension of the theory to cover stochastic equilibrium models.

Briefly, there are two key points: (i) once we fix n and a sample point ω, $f_n(\omega, x)$ becomes a deterministic function of x; (ii) IPA — when it applies — gives *exact* gradients of the f_n. With these observations, very powerful methods of constrained and unconstrained deterministic optimization become available for use on the f_n. In case of both smooth functions and nonsmooth convex functions, we can use superlinearly convergent optimization techniques like the BFGS algorithm –or a variant for constrained problems– (for smooth functions) or the Bundle-Trust (BT) method (for nonsmooth convex functions). These deterministic techniques can explicitly handle several constraints without difficulty and this makes the method attractive to use in the solution of problems with many variables and constraints. It is also worth emphasizing that the availability of powerful deterministic optimization methods enables sample-path optimization to overcome some important difficulties experienced by stochastic approximation methods type methods, Robbins and Monro (1951). These difficulties include sensitivity to the choice of predetermined step sizes and difficulty in enforcing feasibility of iterates, in case of constrained optimization.

4 OUR SOLUTION APPROACH

As a summary, the solution methodology we use consists of fixing a long sample path and optimizing the resulting function using the most powerful deterministic optimization algorithms available to us. Most of these techniques require that function evaluations at a given parameter setting be supplemented by sensitivity information. In our computations, we use infinitesimal perturbation analysis (IPA) to compute sensitivity estimates in the following way.

For a fixed sample path, the long-run average cost per unit time could be defined as $J(z) = C(z,T)/T$, where T is the prespecified time that the controller will be simulated and $C(z,T)$ is the cumulative cost incurred by the system over time T when $z = (z_1, \ldots, z_K)$ is the vector of hedging points. In the following we use $\frac{d^+(\cdot)}{dz_j}$ to denote the one-sided directional derivative of $(\cdot)$ at the point z with respect to the jth unit vector in $\mathbf{R}^K$. If $(\cdot)$ is differentiable at z then $\frac{d^+(\cdot)}{dz_j}$ will be the jth component of $\nabla(\cdot)$.

Let $t_0, t_1, \ldots$ be the event occurrence times in a sample path and $\tau_i = t_i - t_{i-1}$ be the time between the $(i-1)st$ and ith events. Without loss of generality we assume $t_0 = 0$. If the nth event is the event that the controller has been working T time units, i.e. $t_n = T$ for some n, we have

$$\frac{d^+J(z)}{dz_j} = \frac{d^+}{dz_j}\left(\frac{C(z,T)}{T}\right) = \frac{d^+C(z,T)}{dz_j}\frac{1}{T} = \frac{d^+C(z,t_n)}{dz_j}\frac{1}{t_n}. \tag{4.1}$$

From (4.1) we see that $J(z)$ has one-sided directional derivatives if and only if $C(z, t_n)$ has them. In Gürkan and Karaesmen (1998), we show that $C(z, t_n)$ has the desired property and derive a recursive expression for $d^+C(z, t_n)/dz_j$, which is a quantity computable from the simulation information generated up time t_n. Therefore using our IPA algorithm, we can compute the exact one-sided directional derivative of $J(z)$ at any z in a single simulation run (i.e. from a single realization of the sample-path).

Derivation of recursions for the directional derivatives and discussion of “similarity” properties (in the IPA literature) needed for justification of some of these recursions are too detailed and too technical to cover here. However, we would like to point out one strength of our analysis as opposed to other gradient estimation approaches; for example Brémaud *et al.* (1997) and Yan *et al.* (1994). Through the GSMP representation, we derive recursive expressions for computing the directional derivatives; once this analysis has been done, one can use the same recursions to deal with systems with several variables. The same set of recursive expressions can be used to compute sensitivity information with respect to two, three, five, or fifteen hedging points associated with fifteen different production rates. For an illustration, we refer the reader to Gürkan (1999) where a similar analysis has been used to compute the directional derivatives of throughput with respect to buffer capacities, in a tandem production line with 50 unreliable machines.

Furthermore, in the past, an important problem in gradient estimation has been the development of estimators with good asymptotic behavior; this has been mainly related to the convergence analysis of stochastic approximation type methods, see for example Kushner and Clark (1978). It is worth emphasizing an additional advantage of sample-path optimization: the convergence analysis of the conceptual method is independent of the asymptotic properties of the gradient estimator. When using sample-path optimization the only requirement (from the practical point of view) is an exact gradient or a directional derivative (whichever is available) of the sample function to be optimized. If IPA is applicable to a problem, then the resulting gradient estimator naturally satisfies this requirement.

A methodology very similar in spirit has been successfully used in Gürkan (1996, 1999), to solve the optimal buffer allocation problem in tandem production lines with up to 50 machines and 18 linear equality/inequality constraints on the buffer capacities. In those numerical studies, sample-path optimization appeared to give very good results on problems that by current standards would be considered very difficult to solve. We take this as an encouraging sign about the effectiveness and robustness of this solution technique for solving challenging stochastic optimization problems, such as the one we are considering.

5 FUTURE WORK AND EXTENSIONS

This paper summarized some of the recent work we have done for computing optimal hedging points of a manufacturing system (with multiple production rates) that is subject to fairly general random disturbances. Some of this work is still at a preliminary stage and requires further investigation. Future work will be along the following lines:

- Run some experiments and explore the numerical performance of the method suggested.
- Both theoretically and numerically, explore qualitative, functional properties of the cost function (e.g. monotonicity, continuity, convexity, etc.) in hedging points.
- If applicable, investigate whether any of these functional properties can be used to prove the convergence of sample-path optimization.
- If applicable, use the computational procedure to gain more insights about the optimal control of this system.
- If sample-path optimization appears to perform well numerically on this problem, consider a very important extension: adding "probabilistic" constraints. The problem we are currently considering is to minimize the expected cost function, possibly subject to a set of deterministic constraints. Instead of this, it would be very interesting to see what happens if we impose a "probabilistic" constraint such as the long-run probability of backlogs should not exceed some predetermined number α. In principle, we should be able to deal with this extension by using ideas similar to those outlined here; but this requires more careful thinking and possibly additional theoretical work related to sample-path optimization.

REFERENCES

Akella, R. and Kumar, P. R. 1986. Optimal control of production rate in a failure prone manufacturing system. *IEEE Transactions on Automatic Control* 31(2): 116-126.

Bielecki, T. and Kumar, P. R. 1988. Optimality of zero-inventory policies for unreliable manufacturing systems. *Operations Research* 36(4): 532-541.

Brémaud, P., Malhamé, R. P., and Massoulié, L. 1997. A manufacturing system with general stationary failure process: stability and IPA of hedging control policies. *IEEE Transactions on Automatic Control* 42(2): 155-170.

Caramanis, M. and Liberopoulos, G. 1992. Perturbation analysis for the design of flexible manufacturing system flow controllers. *Operations Research* 40(6): 1107-1125.

Glasserman, P. 1991. *Gradient Estimation via Perturbation Analysis* (Kluwer, Norwell, MA).

Glynn, P. W. 1987. Likelihood ratio gradient estimation: an overview. In: *Proceedings of the 1987 Winter Simulation Conference*, eds. A. Thesen, H. Grant, and W.D. Kelton (IEEE, Piscataway, New Jersey), 366-375.

Gürkan, G. 1996. Performance optimization in simulation: Sample-path optimization of buffer allocations in tandem lines. Ph. D. Dissertation. Department of Industrial Engineering, University of Wisconsin-Madison, Madison, WI.

Gürkan, G. 1999. Simulation optimization of buffer allocations in production lines with unreliable machines. To appear in *Annals of Operations Research.*

Gürkan, G. and Karaesmen F. 1998. Computation of optimal flow control policies of a manufacturing system with multiple production rates. Working paper.

Gürkan, G., Özge, A.Y., and Robinson, S.M. 1998. Sample-path solutions for simulation optimization problems and stochastic variational inequalities. In: *Interfaces in Computer Science and Operations Research*, ed. D.L. Woodruff (Kluwer, Boston), 169-188.

Haurie, A., L'Ecuyer, P., and van Delft, C. 1994. Convergence of stochastic approximation coupled with perturbation analysis in a class of manufacturing flow control models. *Discrete Event Dynamic System* 4: 87-111.

Hu, J-Q. and Xiang, D. 1993. The queueing equivalence to a manufacturing system with failures. *IEEE Transactions on Automatic Control* 38(3): 499-502.

Kimemia, J. G. and Gerschwin, S. B. 1983. An algorithm for the computer control of production in a flexible manufacturing system. *IIE Transactions* 15: 353-362.

Kushner, H. J. and Clark, D. S. 1978. *Stochastic Approximation Methods for Constrained and Unconstrained Systems* (Springer-Verlag, New York, NY).

Liberopoulos, G. and Caramanis, M. 1994 Infinitesimal perturbation analysis for second derivative estimation and design of manufacturing flow controllers. *Journal of Optimization Theory and Applications* 81(2): 297-327.

Plambeck, E. L., Fu, B.-R., Robinson, S. M., and Suri, R. 1996. Sample-path optimization of convex stochastic functions. *Mathematical Programming* 75: 137-176.

Reiman, M. I. and Weiss, A. 1986. Sensitivity analysis via likelihood ratios. In: *Proceedings of the 1986 Winter Simulation Conference*, eds. Wilson, R.J., Henriksen J.O., and Roberts, S.D., 285-289.

Robbins, H. and Monro, S. 1951. A stochastic approximation method. *Annals of Mathematical Statistics* 22: 400-407.

Robinson, S. M. 1996. Analysis of sample-path optimization. *Mathematics of Operations Research* 21(3): 513-528.

Sharifnia A. 1988. Production control of manufacturing system with multiple machine states. *IEEE Transactions on Automatic Control* 33: 620-625.

Shedler, G. S. 1993. *Regenerative Stochastic Simulation* (Academic Press, San Diego, CA).

Suri, R. and Fu, B.-R. 1994. On using continuous flow lines to model discrete production lines. *Discrete Event Dynamic Systems* 4: 129-169.

Yan, H., Yin, G., and Lou, S. X. C. 1994. Using stochastic optimization to determine threshold values for the control of unreliable manufacturing systems. *Journal of Optimization Theory and Applications* 83(3): 511-539.