Optimal Manufacturing Flow Controllers: Zero-Inventory Policies and Control Switching Sets*

Michael H. Veatch,† *Gordon College*

Michael C. Caramanis,‡ *Boston University*

June 1998

Abstract

A continuous flow control model of a single workstation with multiple failure modes and part types is considered. Although the general model is intractable, properties of the optimal control policy are identified that can be used to help formulate heuristic policies. Control switching sets are described and shown to have a threshold form. For a single machine with two part types, conditions are found under which no inventory is held, analogous to the single part type result of Bielecki and Kumar.

*This research was supported by National Science Foundation grant DDM-9215368.

†Department of Mathematics, Gordon College, Wenham, MA 01984, (978) 927-2300, FAX (978) 524-3747, veatch@gordonc.edu

‡Department of Manufacturing Engineering, Boston University, Boston, MA 02215, (617) 353-3247, FAX (617) 353-5548, mcc@buenga.bu.edu

1 Introduction

Kimemia and Gershwin [12] proposed a flow control model of a flexible manufacturing system that has been analyzed extensively. They decompose the control decisions in a factory according to the time frames on which they occur. The model considered here focuses on an intermediate time frame, perhaps one or more shifts. In this time frame, demand is treated as constant. The goal is to find real-time production rates to track the demand in response to uncontrollable events. Costs are incurred when there is a production surplus or shortage. When there is a shortage, unmet demand is backlogged until it can be satisfied. The uncertainty of manufacturing environments is captured through Markov changes of state that represent events such as machine failures. A single workstation, consisting of one or more machines, processes several part types. Setup times or costs are not considered.

Because the production of individual parts occurs on a much faster time scale, these events are placed in a lower level in the hierarchy and not considered in this model. Instead, production is modeled as a continuous flow–hence the term *fluid model*. Unlike a two-moment diffusion model, the time to manufacture a part (or, equivalently, the number produced over a time interval) is deterministic for a given machine capacity state. Continuous models can be justified as the limiting high-volume case as buffer sizes, safety stocks or other part counts increase. Even for small part counts (medium- or low-volume manufactuing) continuous models are often close to discrete part models.

Practical algorithms have been developed to compute a control policy for this model using infinitesimal perturbation analysis and some insights have been gained into the structure of good policies [3, 13]. In a slightly more restrictive setting, Sethi et al. [16] prove the existence and convexity of the differential cost function. However, little is known about the structure of the *optimal* policy. This paper investigates that structure. The control switching sets described in [13] are extended to the general problem, conditions are stated under which there are hedging points, and the optimal policy is shown to have a threshold form using the theory of monotone control. These characteristics should prove useful in guiding the development of heuristic policies. As in [16] and our related paper [21], a formulation of the Hamilton-Jacobi-Bellmen (HJB) equations using directional derivatives is employed, placing our results on a more rigorous basis that is not qualified by the assumption of differentiability. For a single machine with two part types, necessary and sufficeint conditions are found under which unreliability does not affect the optimal policy, and no inventory is held. This result extends the single part type result of Bielecki and Kumar. Many of our proofs rely on stochastic coupling, using a combination of sample path and probabilistic arguments to compare two policies.

Understanding of the optimal policy for this problem began with Rishel's [15] results for a terminal reward control problem with Markov changes of state. Single part type problems were solved by Akella and Kumar [1], Bielicki and Kumar [2], and Sharifnia [17] and are summarized in [6]. The optimal policy is a hedging point policy. Srivatsan and Dallery [19, 18] partially characterize the optimal policy for a single machine and two part types. In particular, they show that, for the case of linear surplus and shortage costs, when there is a shortage of both part types a "min $c\mu$" rule applies. They also show that once the high priority shortage is eliminated, a surplus (possibly zero) of this part type is created, then maintained at a constant while the other shortage is eliminated. Finally, the surplus of both types is increased to a hedging point and held constant. Describing the policy by a switching curve between regions in the surplus/shortage space where a part type is produced, the switching curve is vertical in the fourth quadrant, then arbitrary in the first quadrant.

A variant of this problem with setup times is treated by Gershwin, Akella and Choong [7] where the differential cost is approximated by a quadratic. Connelly [4] solves a reliable machine problem with setups. Approximate policies are developed by Veatch and Wein [23, 24] for a related queueing control problem where randomness is incorporated via production times and demand. Queueing methods have also been used by Hu and Xiang [9, 10] and Fu and Hu [5] to analyze variants of the single part type problem with non-exponentially distributed up or down times.

The rest of the paper is organized as follows. Section 2 formulates the problem mathematically, including the HJB equations, and states results from [21] that will be used. The optimal policy is characterized by control switching sets in Section 3, hedging points in Section 4, and threshold form in Section 5. Conditions for optimality of the zero-inventory policy are found in Section 6.

2 The Flow Control Model

We consider the flow control model of Liberopoulos and Caramanis [13], which generalizes the multiple unreliable machine model of [3]. The system state is $(x(t), \alpha(t))$, where $x = (x_1, \dots, x_n)$, x_i is the continuous production surplus of part type i, and α is the discrete machine state. When $x_i(t) > 0$ there is a surplus and when $x_i(t) < 0$ there is a shortage and demand is backlogged. The machine state is governed by a continuous time, irreducible Markov chain on a finite state space $\mathcal{E}$. Let $Q = [q_{\alpha\beta}]$, $\alpha, \beta \in \mathcal{E}$ be the generator, i.e., $q_{\alpha\beta}$ is the transition rate from state α to state β and $q_{\alpha\alpha} = -\sum_{\beta\neq\alpha} q_{\alpha\beta}$. We assume that Q is irreducible and let $\{\pi_\alpha\}$ denote its stationary distribution. Demand occurs at a constant rate d and production occurs at the controllable rate $u(t)$, resulting in the dynamics

$$\dot{x}(t) = u(t) - d. \qquad (1)$$

To simplify notation, production constraints will be stated in terms of velocities $v(t) = \dot{x}(t)$. Many of our results hold for very general production constraints, defined by a convex polyhedron $\mathcal{V}_\alpha$ of feasible velocities $v = \dot{x}$ in state α. Several special cases are of particular interest.

A single machine with maximum production rate μ_i for type i has production constraint $\sum_i u_i/\mu_i \leq 1$. The feasible velocities are

$$\mathcal{V}_1 = \{v : \sum_{i=1}^{n}(v_i + d_i)/\mu_i \leq 1,\ v_i \geq -d_i\} \tag{2}$$

in the "up" state and $\mathcal{V}_0 = \{-d\}$ in the "failed" state . Next consider a two part type system with two machines, the first more efficient at producing type 1 parts and the second more efficient for type 2 parts. The feasible velocity set when both machines are up is shown in Figure 1. The upper right vertex corresponds to producing type 1 with the first machine and type 2 with the second. When either machine is down $\mathcal{V}_\alpha$ is a triangular region, as in (2), and when both machines are down it is the point $-d$. Thirdly, consider an *externally controlled resource:* Setup changes occur randomly and exogenously because the machine is a shared resource, with setups controlled by another shop. In state i only part type i can be produced:

$$\mathcal{V}_i = \{v : -d_i \leq v_i \leq \mu_i - d_i;\ v_j = -d_i, j \neq i\}, \tag{3}$$

$i = 1, \ldots, n$ and $\mathcal{V}_0 = \{-d\}$. The framework can also be used to model piece-wise constant demand that is regulated by a Markov chain.

Cost is incurred at the rate $g(x)$, which is assumed to be convex and additive with a unique minimum at $x = 0$. We also assume that g is polynomially bounded: There are constants C and κ such that

$$g(x) \leq C\left(1 + \sum_{i=1}^{n} |x_i|^\kappa\right) \tag{4}$$

for all x. The objective is to minimize long-run average cost. A control policy is the process $\{v(t) : t \geq 0\}$. In a slight abuse of notation, we will refer to the policy $v(\cdot)$ as v. Policy v is feasible if $v(t) \in \mathcal{V}_\alpha(t)$ for all $t \geq 0$, and admissible if it is feasible and nonanticipating. We state the control problem over the class V_M of stationary feedback policies, i.e., admissible policies v such that $v(t) = v(x(t), \alpha(t))$. However, because $\alpha(\cdot)$ is memoryless and we are considering long-run average cost, methods such as Tsitsiklis [20] can be used to show that it is equivalent to optimize over all admissible policies. The control problem is

$$\begin{aligned} \min_{v} \quad & \limsup_{T\to\infty} \frac{1}{T} E_{x,\alpha} \int_0^T g(x(t))dt && (5) \\ \text{s.t.} \quad & \dot{x}(t) = v(t) && (6) \\ & v \in V_M, && (7) \end{aligned}$$

where $E_{x,\alpha}$ is expectation conditioned on $x(0) = x$ and $\alpha(0) = \alpha$.

We will use the following cost functionals. Denote the cost of policy v in the interval $[0, T]$ by

$$J^v(x, \alpha; T) = E_{x,\alpha} \int_0^T g(x(t))dt, \tag{8}$$

the long-run average cost of policy v by

$$J^v = \limsup_{T\to\infty} J^v(x, \alpha; T)/T, \tag{9}$$

the optimal long-run average cost in (5-7) by J^*, the differential cost, or value function, for policy v by

$$W^v(x, \alpha) = \lim_{T\to\infty} J^v(x, \alpha; T) - TJ^v, \tag{10}$$

and the differential cost for the optimal policy by $W(x, \alpha)$ when these limits exist. The limits (9) and(10) may not exist for all policies.

The usual formulation of the HJB equations for this problem [13] assumes that W is continuously differentiable, which has not been shown. To avoid this difficulty, we use *one-sided directional derivatives*, which are shown to exist in [16] and [21]. Adopt the convention that

$$D_v f(x) = \lim_{h\downarrow 0} \frac{f(x + hv) - f(x)}{h}, \tag{11}$$

a one-sided "directional derivative" with v not normalized. If f is differentiable, then

$D_v f(x) = \nabla f(x) \cdot v$. The HJB equations are

$$J^* = g(x) + \min_{v \in \mathcal{V}_\alpha} D_v W(x, \alpha) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} W(x, \beta) \tag{12}$$

for all $\alpha \in \mathcal{E}$. An informal derivation of (12), such as [3] or [6, Section 8.8], can be adapted to use only directional derivatives of W.

Now we return to the question of whether J^v and W^v exist. A policy v is *stable* if

$$\lim_{t \to \infty} \frac{E_{x,\alpha} |x_i(t)|^{\kappa+1}}{t} = 0$$

for $i = 1, \ldots, n$. Here the dependency of x on v is suppressed and κ bounds the order of the cost rate g. The following facts have been established for similar models and can be demonstrated for our model using the same methods.

(i) If x has a stationary distribution under v, then v is stable.

(ii) If v is stable, then J^v exists (see [11]).

(iii) If there is a recurrent state (x, α) under v, then W^v exists (see [20]).

Define the expected capacity set [3]

$$\overline{\mathcal{V}} = \{v : v = \sum_{\alpha \in \mathcal{E}} \pi_\alpha v^\alpha,\ v^\alpha \in \mathcal{V}_\alpha \text{ for all } \alpha\}.$$

In order to be stable, a system must have excess capacity: There is a $v \in \overline{\mathcal{V}}$ such that $v > 0$. Single-machine [19, Section 2] and multiple-machine [3] models have stable policies if and only if there is excess capacity. However, excess capacity is not sufficient in our general model. For example, systems that can only produce part types in a fixed ratio may not be stable. We do not have a complete specification of which systems are stable; nevertheless, we restrict our attention to stable systems.

Up to this point $\mathcal{V}_\alpha$ is very general; in particular, surpluses and shortages are symmetric, so that a hedging point could occur at $x_i < 0$. The following assumptions will be used as necessary to consider more restrictive (and realistic) cases.

No forced overproduction: For any $v \in \mathcal{V}_\alpha$ with $v_i > 0$, the velocity v' differing from v only in the ith component being zero is also in $\mathcal{V}_\alpha$.

Failed state: For some $\alpha \in \mathcal{E}$, say $\alpha = 0$, $\mathcal{V}_0 = \{-d\}$, where $d > 0$.

We will use the following theorems from [21].

Theorem 1 *$W(x, \alpha)$ is convex in x. If there is a failed state, then $W(x, \alpha)$ is strictly convex on $\mathbf{R}^n_+$.*

Theorem 2 *For machine states α with 0 an interior point of $\mathcal{V}_\alpha$, $W(\cdot, \alpha)$ is differentiable at the hedging point z^α.*

The proof of Theorem 2 also applies to W^v where v is the zero-inventory policy.

3 Control Switching Sets

Special cases of this problem have exhibited a certain structure of the optimal policy which Liberopoulos and Caramanis [3], [13] have called *control switching sets* (CSSs). In this section we adapt their CSS structure to the general problem.

For the single-machine problem of [13], CSSs partition the x-space into regions in which different extreme points of $\mathcal{V}_\alpha$ are optimal. Boundaries where two or more CSSs intersect form lower dimensional sets where multiple extreme points of $\mathcal{V}_\alpha$ are optimal. In intervals of constant $\alpha(t)$, the optimal x-trajectory moves deterministically, instantly passing through *deflective* region boundaries when one optimal extreme point is replaced by another and remaining in *attractive* region boundaries when an extreme point is added to the optimal set. Figure 2 shows a trajectory that begins in the CSS where v^j is optimal, reaches the attractive CSS where v^i and v^j are optimal, then stops at the hedging point (the notation in Figure 2 will be defined later). If 0 is an interior point of $\mathcal{V}_\alpha$, the deterministic trajectory terminates at a *hedging point*, which is the minimum of $W(x, \alpha)$. If $0 \notin \mathcal{V}_\alpha$, $x(t)$ continues to drift until α changes. If 0 is on the boundary of $\mathcal{V}_\alpha$, the deterministic trajectory may reach the hedging point or may terminate elsewhere, depending on the initial x.

We make two adjustments to their structure: (i) because we do not assume that W is continuously differentiable, (12) may not have extreme point solutions for all x and (ii) a system can be stable without having a hedging point. Let $\mathcal{V}^*(x)$ be the set of controls v that achieve the minimum in (5-7) (the dependency on α is suppressed) and Conv denote the convex hull. For any set $\{v^1, \dots, v^l\}$ of extreme points of $\mathcal{V}_\alpha$, the CSS is defined as

$$X(v^1, \dots, v^l) = \{x : \{v^1, \dots, v^l\} \text{ is minimal s.t. } \mathcal{V}^*(x) \subseteq \text{Conv}(v^1, \dots, v^l)\}. \quad (13)$$

If W is not differentiable at x, then $\mathcal{V}^*(x)$ may not contain an extreme point of $\mathcal{V}_\alpha$; instead, it may be in the interior of some face $\text{Conv}(v^1, \dots, v^l)$ of $\mathcal{V}_\alpha$.

An attractive CSS, in the sense of [13], has the property that it is possible to remain in the CSS for a nonzero amount of time using a control in $\text{Conv}(v^1, \dots, v^l)$. Roughly, an attractive CSS extends in a direction that is in the convex cone of the vectors $v^1, \dots, v^l$. A CSS is also attractive if 0 is a convex combination of the vectors $v^1, \dots, v^l$. It is conceivable that part of a CSS has the attractive property and part does not. If a set of points in the CSS with the same dimension as the whole CSS does not have the attractive property, we say the CSS is deflective.

In CSSs with more than one optimal control, the control can be chosen so as to avoid chattering, as discussed in [13]. Typically, a CSS of l extreme points has dimension $n-l+1$ or is the empty set. In this case, CSSs where only one extreme point is optimal divide $\mathbf{R}^n$ into regions, and tie-breaking is needed only on their boundaries. However, the single-machine problem with symmetric part types illustrates that CSSs can have larger dimensions. In this problem, the CSS for the n extreme points corresponding to producing each part type includes all $x < 0$ and has dimension n.

4 Hedging Points

As noted above, a system can be stable without having a hedging point. The externally controlled resource example (3) does not have a hedging point because, for all α, $0 \notin \mathcal{V}_\alpha$. The part types are independent and decompose into single part type problems with multiple failure states, all of which may be stable. Clearly, to remain at a hedging point it must be the case that $0 \in \mathcal{V}_\alpha$. Using some natural additional conditions, we show that (stable) systems have hedging points.

Lemma 1 *If there is no forced overproduction, then $W(x, \alpha)$ is decreasing in x_i when $x_i < 0$.*

Proof. Consider the coupled process $x^0(t)$, $x^1(t)$ with initial states x^0 and x^1 sharing the same machine state process $\alpha(t)$. Let $x^1 = x^0 + \Delta e_i$, where $x_i^1 \le 0$, $\Delta > 0$, and e_i is the unit vector with ith component equal to one. Use the optimal policy for $x^0(t)$ and the same (nonstationary) control $v^1(t) = v^*(x^0(t), \alpha(t))$ for $x^1(t)$ except that the boundary $x_i^1 \le 0$ is enforced until they merge. The assumption of no forced overproduction makes this policy feasible. Then $0 \ge x_i^1(t) > x_i^0(t)$ and $g(x^1(t)) < g(x^0(t))$ prior to merging. It follows that $J^{v^*}(x^1, \alpha; T) \le J^{v^1}(x^1, \alpha; T) < J^{v^*}(x^0, \alpha; T)$ and $W(x^1, \alpha) < W(x^0, \alpha)$. $\square$

Lemma 2 *If there is no forced overproduction and there is a failed state, then $W(x, \alpha)$ has a unique minimum.*

Proof. Since g is convex with a unique minimum, it grows at least linearly in $\|x\|$. Stability, i.e., the existence of J^*, implies that $x(t)$ enters the region near the origin where $g(x) \leq J^*$ with probability one. Then, using the fact that $\dot{x}$ is bounded, one can establish that $W(x, \alpha)$ grows at least quadratically in $\|x\|$ (we omit the details). Hence, $W(x, \alpha)$ has a minimum. Lemma 1 implies that the minimum is in $\mathbf{R}^n_+$. Strict convexity on $\mathbf{R}^n_+$ (Theorem 1) guarantees that the minimum is unique. □

Let z^α minimize $W(\cdot, \alpha)$.

Theorem 3 *If there is no forced overproduction, there is a failed state, and* $0 \in \mathcal{V}_\alpha$*, then* $z^\alpha \geq 0$ *is the hedging point for machine state* α*.*

The theorem follows immediately from Lemma 2 and (12). Under the optimal policy, after $x(t)$ reaches the hedging point it will remain there until the machine state changes. This makes the hedging point an attractive CSS.

5 Monotone Switching

In this section, the optimal policy is shown to have a threshold, or monotone, form. This restricts the boundaries of each control region to lie within certain directions determined by the geometry of $\mathcal{V}_\alpha$. Our approach is taken from the theory of monotone control for discrete state dynamic programs, described in Veatch and Wein [22] and Glasserman and Yao [8].

Theorem 4 *If* v^i *is an optimal control in state* (x, α)*, then moving in* x*-space in the direction* $v^i - v^j$ *may change the optimal control to* v^j*, but moving in the opposite direction will not.*

Proof. The control v^i is prefered over v^j in (12) when $D_{v^i}W(x,\alpha) \leq D_{v^j}W(x,\alpha)$. Fix x and consider the line $x + s(v^i - v^j)$ with parameter s (Figure 2). We will show that this preference changes monotonically as s increases. Let

$$\begin{aligned} f_i(s) &= D_{v^i}W(x + s(v^i - v^j)), \\ f_{i-j}(s) &= D_{v^i - v^j}W(x + s(v^i - v^j)), \quad \text{and} \\ s_{ij} &= \min\{s : f_{i-j}(s) \geq 0\}. \end{aligned}$$

Convexity of W (hence, of f) implies that f_{i-j} is increasing in s. It also implies that $f_i(s) - f_j(s) \leq f_{i-j}(s)$, since the linear approximation of W in the $v^i - v^j$ direction cannot lie below the plane established by the v^i and v^j directions (if W is continuously differentiable, then equality holds). For $s < s_{ij}$,

$$f_i(s) - f_j(s) \leq f_{i-j}(s) < 0$$

and v^i is prefered. For $s > s_{ij}$,

$$f_j(s) - f_i(s) \leq f_{j-i}(s) \leq -f_{i-j}(s_{ij}) \leq 0$$

and v^j is prefered. Thus, the preference switches monotonically from v^i to v^j at the point where $s = s_{ij} > 0$, and the theorem follows. □

Combining the switching points s_{ij} for all initial x gives a *switching surface* where the preference changes from v^i to v^j. Of course, it may be that some other control is optimal on all or part of the switching surface. If v^i and v^j are adjacent extreme points of $\mathcal{V}_\alpha$ and their CSSs $X(v^i)$ and $X(v^j)$ are nonempty, then one would expect $X(v^i, v^j)$ to be nonempty and to contain the part of the switching surface where no other controls are optimal. If v^i and v^j are not adjacent, one would not expect the optimal control to switch directly from v^i to v^j, and $X(v^i, v^j)$ would be the empty set.

Now we apply the theorem to a single-machine problem. Each nonsingleton $\mathcal{V}_\alpha$ has the form (2) and the extreme points $v^0 = -d$ and $v^i = (-d_1, \ldots, \mu_i - d_i, \ldots, -d_n)$, $i = 1, \ldots, n$. These controls represent idling and producing type i, respectively. Theorem 4 states that moving in the $v^i - v^0$ direction (increasing x_i) may change the optimal control from producing type i to idling, but not vice-versa. We can use monotonicity to characterize the *idleness region* $X(v^0)$ in which idling is optimal. If $X(v^0)$ is nonempty, and assuming ties occur only on sets $X(v^0, v^i)$ of lower dimension, define the switching surfaces

$$s_{0i}(x_{(i)}) = \min\{x_i : D_{e_i} W(x, \alpha) \geq 0\},$$

where $x_{(i)}$ denotes x with the ith component removed. These surfaces bound the idleness region, i.e., letting $\overline{X}$ denote the closure of X,

$$\overline{X(v^0)} = \{x : x_i \geq s_{0i}(x_{(i)})\}.$$

This rather mild characterization agrees with numerical experience. That experience also suggests the stronger condition that s_{0i} is nonincreasing in x_j, $j \neq i$ (see [22] for related discussion).

Similarly, moving in the direction $v^i - v^j = \mu_i e_i - \mu_j e_j$ may change the optimal control from producing type i to producing type j, but not vice versa. Note that moving in this direction keeps the workload $\sum_i x_i/\mu_i$ constant, increasing x_i and decreasing x_j. A v^i, v^j switching surface can be defined that intersects each line of constant workload in at most one point. The region $X(v^i)$ in which producing type i is optimal, then, is the region on the side with decreasing x_i of each switching surface, $j \neq i$, excluding the idleness region. Again, monotonicity places mild restrictions on the switching surfaces. A stronger condition, that on the switching surface between v^i and v^j, x_i is nondecreasing in x_j, seems likely.

6 A Zero-Inventory Policy

In this section we focus on the two part type single-machine problem and find conditions for a zero-inventory policy to be optimal. Srivatsan [18] uses formulas for W^v, where they are known, to find partial conditions. We evaluate the effect of three hedging point perturbations using stochastic coupling and obtain simple necessary and sufficient conditions for optimality. To illustrate the power of this technique, we also use it to locate the vertical switching curve of [19] Theorem 4.

Specializing (2) to two part types, the velocity constraint in the "up" state ($\alpha = 1$) is $(v_1 + d_1)/\mu_1 + (v_2 + d_2)/\mu_2 \leq 1$, where μ_i and d_i are the maximum production rate and demand for type i. In the "down" state ($\alpha = 0$), $v = -d$. Shortage and surplus costs are linear, $g(x) = \sum_i (g_i^- x_i^- + g_i^+ x_i^+)$. Assume that $g_1^- \mu_1 \geq g_2^- \mu_2$, so that it is optimal to produce type 1, $v = (\mu_1 - d_1, -d_2)$, when there is a shortage of both types. Let p be the failure rate and r the repair rate. The stability condition for this system is

$$\rho \equiv \frac{d_1}{\mu_1} + \frac{d_2}{\mu_2} < \frac{r}{r+p},$$

where ρ is the utilization of a reliable system and the right side is $P(\alpha = 1)$. Consider a policy partially characterized by

$$v(x,1) = \begin{cases} (\mu_1 - d_1, -d_2) & \text{if } x_1 < z_1^m \text{ and } x_2 < 0 \\ (0, \mu_2(1 - d_1/\mu_1) - d_2) & \text{if } x_1 = z_1^m \text{ and } x_2 < 0 \\ (\mu_1(1 - d_2/\mu_2) - d_1, 0) & \text{if } z_1^m \leq x_1 \leq \overline{z}_1 \text{ and } x_2 = 0. \end{cases} \quad (14)$$

In this policy, the switching curve $X(v^1, v^2)$ between producing type 1 and type 2 is vertical for $x_2 < 0$, then at $x_2 = 0$ it has a horizontal segment (Figure 3b). Srivatsan proves that optimal policies must have this characteristic. To find z_1^m, consider the incremental cost of a small shift in the vertical switching curve. Let x^1 use the policy above and x^2 use a policy with the vertical switching curve shifted to $z_1^m + \epsilon < \overline{z}_1$. Both processes have the same machine states. Typical trajectories are shown in Figure 4.

From the initial state $x^i(0) = (z_1^m, x_2)$ and $\alpha(0) = 1$, x^1 will use the first control in (14) while x^2 uses the second. If no failure occurs, these controls are used until x^2 reaches its switching curve. At this point

$$x^2(\tau_1) - x^1(\tau_1) = \epsilon(1, -\mu_1/\mu_2) \tag{15}$$

and their relative position remains constant until τ_2, when $x^2 = (z_1^m, 0)$. Then x_2 switches to the third control in (14) and, if no failure occurs, they merge at $x^1(\tau_3) = x^2(\tau_3) = (z_1^m + \epsilon, 0)$. From here they will remain merged, $x^1 = x^2$, until they again reach the vertical switching curve. Hence, τ_3 is a renewal point of the coupled process and minimizing cost until renewal is equivalent to minimizing long-run average cost.

For small ϵ, the relative costs and transitions in the intervals $[0, \tau_1]$ and $[\tau_2, \tau_3]$ can be neglected. Hence, we consider only the relative position in (15), with incremental cost rate

$$\Delta g = g(x^2) - g(x^1) = \begin{cases} \epsilon(-g_1^- + g_2^- \mu_2/\mu_1) & \text{if } x_1^2 \leq 0 \\ \epsilon(g_1^+ + g_2^- \mu_2/\mu_1) & \text{if } x_1^1 \geq 0. \end{cases} \tag{16}$$

The expected relative cost incurred when $x_1^1 < 0 < x_1^2$ is also $o(\epsilon)$ and can be neglected. Thus, for small ϵ, the relative cost until renewal is the cost of a single part type 1 system under a hedging point policy with the modified shortage and surplus costs given in (16). To minimize this cost, we choose z_1^m as the optimal hedging point for this single part type problem, which is Srivatsan's result.

Now we use a similar argument to obtain necessary and sufficient conditions for optimality of a zero-inventory policy, defined by

$$v(x, 1) = \begin{cases} (\mu_1 - d_1, -d_2) & \text{if } x_1 < 0 \text{ and } x_2 \leq 0 \\ (0, \mu_2(1 - d_1/\mu_1) - d_2) & \text{if } x_1 = 0 \text{ and } x_2 < 0 \\ (0, 0) & \text{if } x_1 = 0 \text{ and } x_2 = 0 \end{cases} \tag{17}$$

in its recurrent states (Figure 3a). From the initial state $x = (0, 0)$ and $\alpha = 1$, the trajectory consists of (1) remaining in the initial state until a failure occurs, (2) an initial shortage of part type 1, (3) one or more intervals moving up the x_2 axis, and (4) possible subsequent shortages of part type 1. Figure 5 shows a typical trajectory. Let p_1, p_2, p_3, and p_4 be the probabilities that a randomly selected time between renewals lies in these four respective parts of the trajectory. To find these probabilities, first consider a single part type problem . Dropping the part type subscript, the probability of being at the hedging point z satisfies the flow balance equation

$$(\mu - d)[1 - P(X = z) - P(\alpha = 0)] = dP(\alpha = 0),$$

resulting in

$$P(X = z) = 1 - \frac{p}{r + p}\left(\frac{1}{1 - d/\mu}\right).$$

For our system, $p_1 = P(X = z)$ is the same as the probability that a single part type system with utilization ρ is at its hedging point, and we have

$$p_1 = 1 - \frac{p}{r + p}\left(\frac{1}{1 - \rho}\right). \tag{18}$$

Furthermore, $p_1 + p_3 = P(X_1 = 0)$ is the hedging probability for part type 1 in isolation, and

$$p_1 + p_3 = 1 - \frac{p}{r + p}\left(\frac{1}{1 - d_1/\mu_1}\right). \tag{19}$$

To find p_2, let $N_i(t)$ be the number of failures in $(0, t]$ occuring while in part $i =$ 1,2,3, or 4 of the trajectory and $T_i(t)$ be the time in $(0, t]$ spent in part i. Since failures occur according to a Poisson process, a weak law of large numbers argument can be used to show that the proportion of failures on each part approaches the fraction of time in each part with probability one as $t \to \infty$. But the fraction of time in part i also approaches p_i with probability 1. For example,

$$\lim_{t\to\infty} \frac{N_1(t)}{N_1(t) + N_3(t)} = \lim_{t\to\infty} \frac{T_1(t)}{T_1(t) + T_3(t)} = \frac{p_1}{p_1 + p_3} \tag{20}$$

with probability one. Each failure counted in N_1 or N_3 initiates a part type 1 shortage, with failures in part 1 of the trajectory entering part 2 and failures in part 3 of the trajectory entering part 4. These relationships imply that

$$\lim_{t\to\infty} \frac{T_2(t)}{T_2(t)+T_4(t)} = \lim_{t\to\infty} \frac{N_1(t)}{N_1(t)+N_3(t)} \tag{21}$$

with probability one. Combining (20) and a similar equation with (21), we conclude that

$$p_2 = \frac{p_1}{p_1+p_3}(p_2+p_4) = \frac{p_1}{p_1+p_3}(1-p_1-p_3). \tag{22}$$

Finally, $p_4 = 1 - p_1 - p_2 - p_3$. We will use the notation $\gamma = p_1 = P(X = z)$ and $\gamma_1 = p_1 + p_3 = P(X_1 = z_1)$.

Three perturbations of the policy will be considered: moving the hedging point up, shifting the switching curve to the right, and adding a horizontal segment to the switching curve by moving the hedging point to the right. We claim that the zero-inventory policy is optimal if and only if none of these perturbations reduces cost. Clearly, if a perturbation decreases cost the policy is not optimal. Long-run average cost can be shown to have the following properties:

1. J is a convex function of horizontal shifts of the policy.
2. Consider the class of policies that use the optimal switching curve but various hedging points along this curve. J decreases monotonically as the hedging point moves toward the optimal point.
3. For the zero-inventory policy, the change in J caused by extending the switching curve up or to the right is *not decreased* by shifting the switching curve to the right.

Suppose some other policy v' of the form (14) is optimal and none of the three perturbations of the zero-inventory policy decrease cost. The v' policy (Figure 3b) can be obtained from the zero-inventory policy (Figure 3a) by shifting its switching curve to the right, then extending it in directions that lie in the first quadrant. By property (1) the horizontal shift increases J. Property (3) implies that extending the switching curve up or to the right still increases J after the shift. Differentiability of W^v at the hedging point (for the zero-inventory policy v) allows us to make the same statement for any direction in the first quadrant. But this contradicts property (2). Hence, the zero-inventory policy must be optimal.

First, consider the incremental cost of shifting the switching curve to the right. We will analyze the expected cost between renewals at the hedging pont, since the time between renewals is unchanged. Define two processes, x^1 using the zero-inventory policy and x^2 using the hedging point $(\epsilon, 0)$ with a vertical switching curve. Both use the same machine states. Let $x^1(0) = (0,0)$, $x^2(0) = (\epsilon, 0)$, and $\alpha(0) = 1$. Neglecting the $o(\epsilon)$ effect of being in the region $0 < x_1^2 < \epsilon$, the relative cost rates are

$$\Delta g = g(x^2) - g(x^1) = \begin{cases} \epsilon g_1^+ & \text{part 1 and 3} \\ -\epsilon g_1^- & \text{part 2 and 4.} \end{cases}$$

The time-averaged incremental cost rate, found by averaging over a renewal period, is

$$\Delta J_{right} = \epsilon[(p_1 + p_3)g_1^+ - (p_2 + p_4)g_1^-]. \tag{23}$$

Similarly, moving the hedging point up to $(0, \epsilon)$ and using initial states $x^1(0) = (0,0)$ and $x^2(0) = (0, \epsilon)$ produces relative cost rates of ϵg_2^+ in part 1 and $-\epsilon g_2^-$ otherwise, and

$$\Delta J_{up} = \epsilon[(p_1 g_2^+ - (1 - p_1)g_2^-]. \tag{24}$$

Thirdly, extending the switching curve horizontally and using the same initial states as the first case, the relative cost rates are

$$\Delta g = \begin{cases} \epsilon g_1^+ & \text{part 1} \\ -\epsilon g_1^- & \text{part 2} \\ -\epsilon \frac{\mu_2}{\mu_1} g_2^- & \text{part 3 and 4.} \end{cases}$$

Note that in this case the relative position $x^2 - x^1$ is initially $(\epsilon, 0)$ but changes to

$(0, \epsilon\mu_2/\mu_1)$ at the end of part 2 of the trajectory. At this point, the first shortage of type 1 ends and both x^1 and x^2 encounter the same switching curve. The time-averaged incremental cost rate is

$$\Delta J_{horiz} = \epsilon \left[(p_1 g_1^+ - p_2 g_1^- - (p_3 + p_4) \frac{\mu_2}{\mu_1} g_2^- \right]. \tag{25}$$

As argued earlier, a necessary and sufficient condition for optimality of the zero-inventory policy is that the incremental costs (23), (24), and (25) are nonnegative. To make these conditions more transparent, rewrite (22) as $p_2 = (\gamma/\gamma_1)(1 - \gamma_1)$, and observe that $p_3 + p_4 = 1 - p_1 - p_2 = (\gamma_1 - \gamma)/\gamma_1$. With these substitutions, the zero-inventory conditions are

$$\gamma_1 \geq \frac{g_1^-}{g_1^+ + g_1^-}, \tag{26}$$

$$\gamma \geq \frac{g_2^-}{g_2^+ + g_2^-}, \text{ and} \tag{27}$$

$$\gamma_1 \geq \frac{g_1^-}{g_1^+ + g_1^-} + \left(\frac{\gamma_1 - \gamma}{\gamma}\right)\left(\frac{\mu_2}{\mu_1}\right)\frac{g_2^-}{g_1^+ + g_1^-}. \tag{28}$$

Since $\gamma < \gamma_1$, the last term in (28) is positive and (28) implies (26); i.e., (26) can be omitted. In terms of the perturbations used, the switching curve need not be shifted to the right–a horizontal extension suffices. This dominance suggests that a vertical switching curve at $x_1 > 0$ cannot be optimal. Conditions (26) and (27) are the same as Bielicki and Kumar's zero-inventory condition for the appropriate single part type problems. Because part type 1 is given static priority, it sees only type 1 parts, while part type 2 in effect sees the combined workload of both types.

Also, (26) and (27) are the *necessary* conditions for zero inventory given by Perkins and Srikant Theorem 2.2. Our *necessary and sufficient* conditions are stronger than their necessary conditions. Our conditions are also weaker than their *sufficient* conditions (Theorem 2.1), which are

$$\gamma \geq \frac{g_1^-}{g_1^+ + g_1^-} \tag{29}$$

and (27). They are weaker because (29) and Perkins and Srikant's assumption that $\mu_1 = \mu_2$ imply (28). The basic reason for our conditions being necessary and sufficient is that we consider a policy that does not give static priority to part type 1, namely, adding a horizontal segment to the switching curve. If $g_1^- \mu_1 = g_2^- \mu_2$ (the parts are tied for priority when both have shortages), then, because the part types can be renumbered, all three sets of conditions are equivalent. One way to view this reduction is to redefine $\gamma_1 \equiv \gamma$.

The difference between the conditions can be significant. Set $\mu_1 = \mu_2 = 10$, $d_1 = 2$, $d_2 = 4$, $g_1^+ = 1$, and $g_2^+ = 2$. If $g_1^- = 7$ and $g_2^- = 6$, the necessary conditions (26) and (27) are satisfied but zero-inventory is not optimal. If $g_1^- = 4$ and $g_2^- = 2$, zero-inventory is optimal but the sufficient condition (29) is not satisfied. For smaller shortage costs, such as $g_1^- = 2.3$ and $g_2^- = 2.1$, the sufficient conditions (26) and (29) are satisfied.

Our conditions indicate, as do Bielicki and Kumar's, that inventory should never be held in systems where surplus costs are large relative to shortage costs, nor in systems where reliability is high and there is significant excess capacity. Condition (28) raises a new issue: If the lower priority part type has shortage costs that are significant relative to the costs of the higher priority part type and uses a significant percentage of capacity, then holding inventory is more likely to be justified.

Acknowledgements

We thank Stan Gershwin and Jim Perkins for their helpful comments and our students Shirish Ekbote, Chang Shu, and Michael Yee for their analysis of when the zero-

inventory conditions hold.

References

[1] R. Akella and Kumar, P. R., "Optimal Control of Production Rate in a Failure Prone Manufacturing System," *IEEE Trans. Automat. Contr.,* Vol. AC-31, pp.116-126, 1986.

[2] T. Bielecki and Kumar, P. R., "Optimality of Zero-Inventory Policies for Unreliable Manufacturing Systems," *Oper. Res.,* Vol. 36, pp. 532-541, 1988.

[3] M. Caramanis and G. Liberopoulos, "Perterbation Analysis for the Design of Flexible Manufacturing System Flow Controllers," *Oper. Res.,* Vol. 40, pp. 1107-1125, 1992.

[4] S. Connelly, "A Real-Time Policy for Performing Setup Changes in a Manufacturing System," Masters thesis, Operations Research Center, MIT, 1992.

[5] M. C. Fu and J. Q. Hu, "On the Relationship of Capacitated Production/Inventory Models to Manufacturing Flow Control Models," *Operations Research Letters,* vol. 18, pp. 15-24, 1995.

[6] S. B. Gershwin, *Manufacturing Systems Engineering,* Englewood Cliffs, New Jersey: Prentice-Hall, Inc., 1994.

[7] S. B. Gershwin, R. Akella, and Y. F. Choong, "Short-Term Production Scheduling of an Automated Manufacturing Facility," *IBM J. Res. Develop.,* Vol. 29, pp. 392-400, 1985.

[8] P. Glasserman and D. D. Yao, "Montone Optimal Control of Permutable GSMPs," *Math. of Oper. Res.,* Vol. 19, pp. 449-476, 1994.

[9] J. Q. Hu and D. Xiang, "A Queueing Equivalence to a Manufacturing System with Failures," *IEEE Trans. Automat. Contr.,* Vol. 38, pp. 499-502, 1993.

[10] J. Q. Hu and D. Xiang, "Optimal Control for Systems with Deterministic Production Cycles," *IEEE Trans. Automat. Contr.,* Vol. 40, No.4, pp.782-786, 1995.

[11] J. Kimemia, "Hierarchical Control of Production in Flexible Manufacturing Systems," Ph.D. dissertation, Laboratory for Information and Decision Sciences, MIT, Cambridge, Massachusetts, 1982.

[12] Kimemia, J., and Gershwin, S. B., "An Algorithm for the Computer Control of a Flexible Manufacturing System," *IIE Transactions,* Vol. 15, pp. 353-362, 1983.

[13] G. Liberopoulos and M. Caramanis, "Dynamics and Design of a Class of Parameterized Manufacturing Flow Controllers," *IEEE Trans. Automat. Contr.,* Vol. 40, No. 6, 1995.

[14] J. R. Perkins and R. Srikant, "Hedging Policies for Failure-Prone Manufacturing Systems: Optimality of JIT and Bounds on Buffer Levels," to appear *IEEE Trans. Automat. Contr.,* July 1998.

[15] R. Rishel, "Dynamic Programming and Minimal Principles for Systems with Jump Markov Disturbances," *SIAM Jour. Contr.,* Vol. 13, pp. 338-371, 1975.

[16] S. P. Sethi, W. Suo, M. I. Taksar, and H. Yan, "Optimal Production Planning in a Multi-Product Stochastic Manufacturing System with Long-Run Average Cost," *Discrete Event Dynamical Systems,* Vol. 8, pp. 37-54, 1998.

[17] A. Sharifnia, "Production Control of a Manufacturing System with Multiple Machine States," *IEEE Trans. Automat. Contr.,* Vol. 33, pp. 620-625, 1988.

[18] N. Srivatsan, *Synthesis of Optimal Policies in Stochastic Manufacturing Systems,* Ph.D. disseration, Operations Research Center, MIT, 1993.

[19] N. Srivatsan and Y. Dallery, "Partial Characterizations of Optimal Hedging Point Policies in Unreliable Two-Part-Type Manufacturing Systems," *Oper. Res.,* Vol. 46, pp. 36-45, 1998.

[20] J. N. Tsitsiklis, “Convexity and Characterization of Optimal Policies in a Dynamic Routing Problem,” *J. Optimiz. Th. Applic.,* Vol. 44, pp. 105-136, 1984.

[21] M. H. Veatch and M. C. Caramanis, “Optimal Average Cost Manufacturing Flow Controllers: Convexity and Differentiability,” to appear *IEEE Trans. Automat. Contr.*

[22] M. H. Veatch and L. M. Wein, “Monotone Control of Queueing Networks,” *Queueing Systems Theory and Applic.,* Vol. 12, pp. 391-408, 1992.

[23] M. H. Veatch and L. M. Wein, “Scheduling a Make-to-Stock Queue: Index Policies and Hedging Points,” *Oper. Res.* Vol. 44, pp. 634-647, 1996.

[24] L. M. Wein, “Dynamic Scheduling of a Multiclass Make-to-Stock Queue,” *Oper. Res.,* Vol. 40, pp. 724-735, 1992.

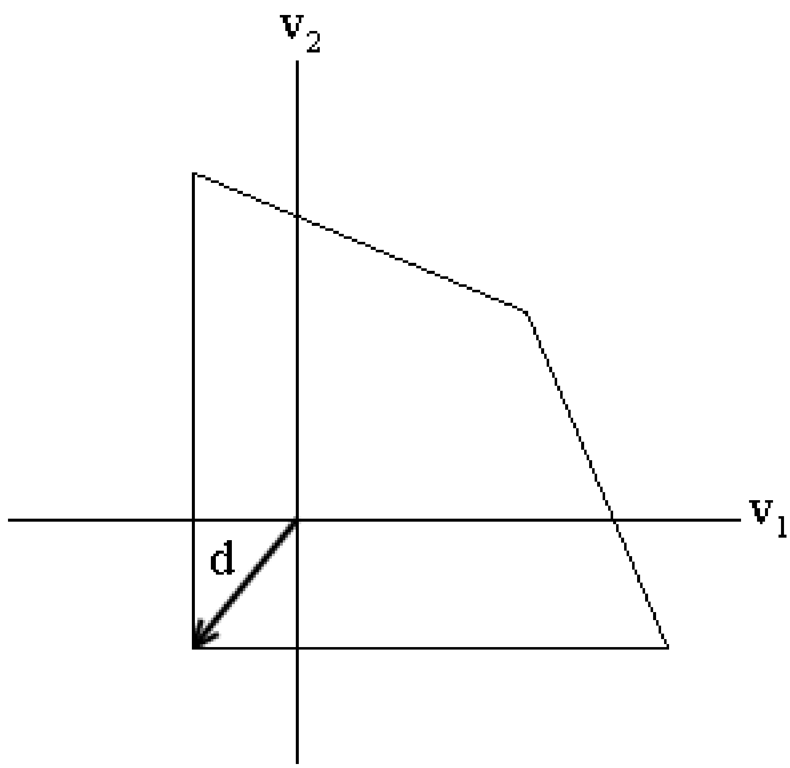

Figure 1: Feasible Velocities for a Two-Machine System.

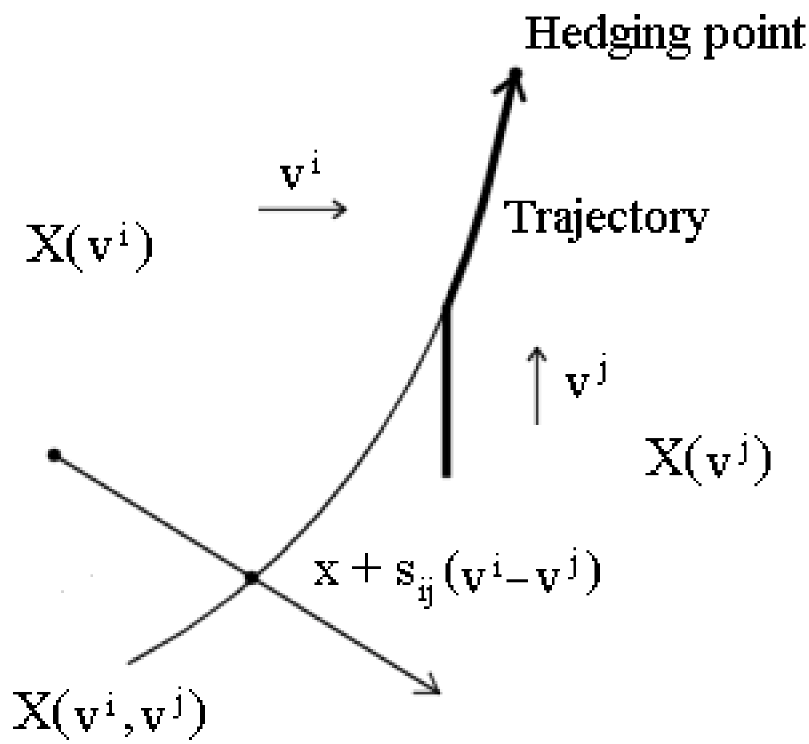

Figure 2: An Attractive Control Switching Set

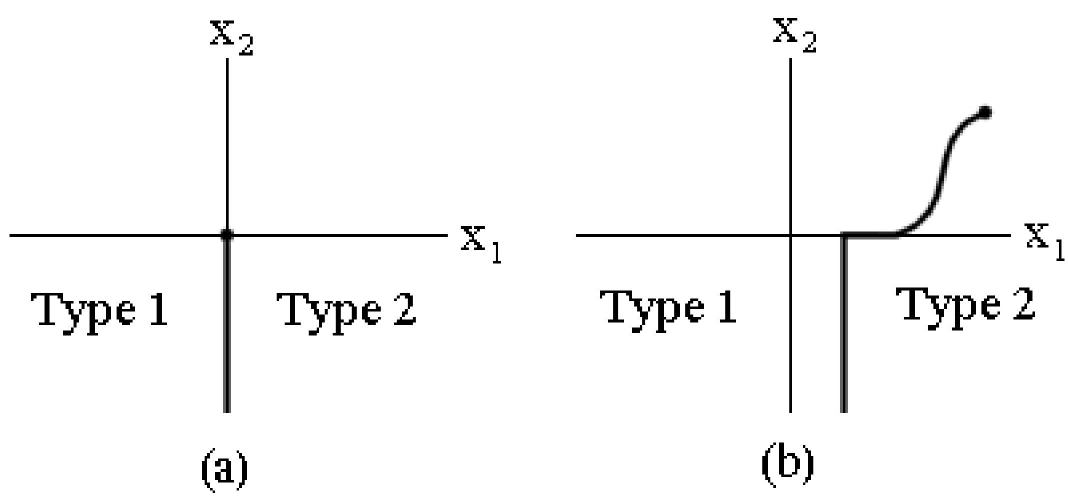

Figure 3: a) The Zero-Inventory Policy. b) A General Optimal Policy.

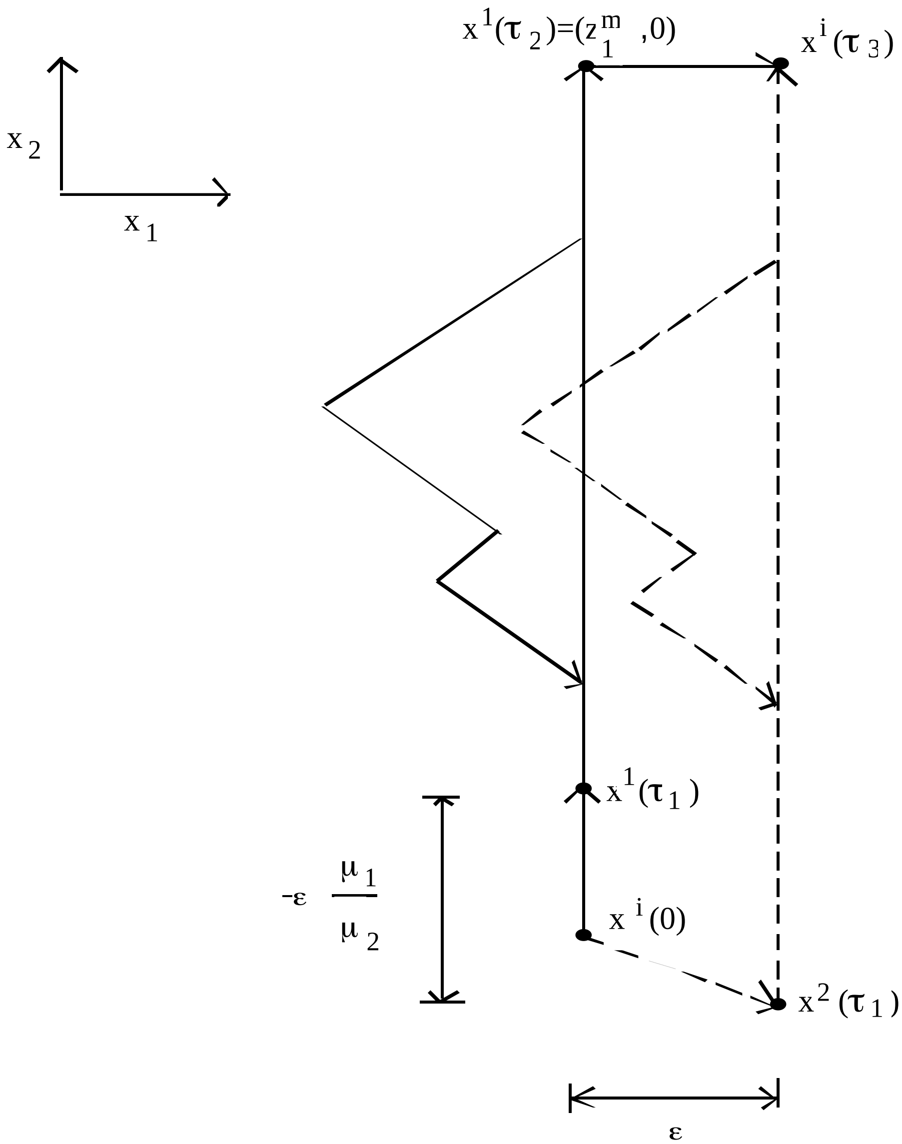

Figure 4: Typical Trajectory for the Coupled Processes x^1 and x^2.

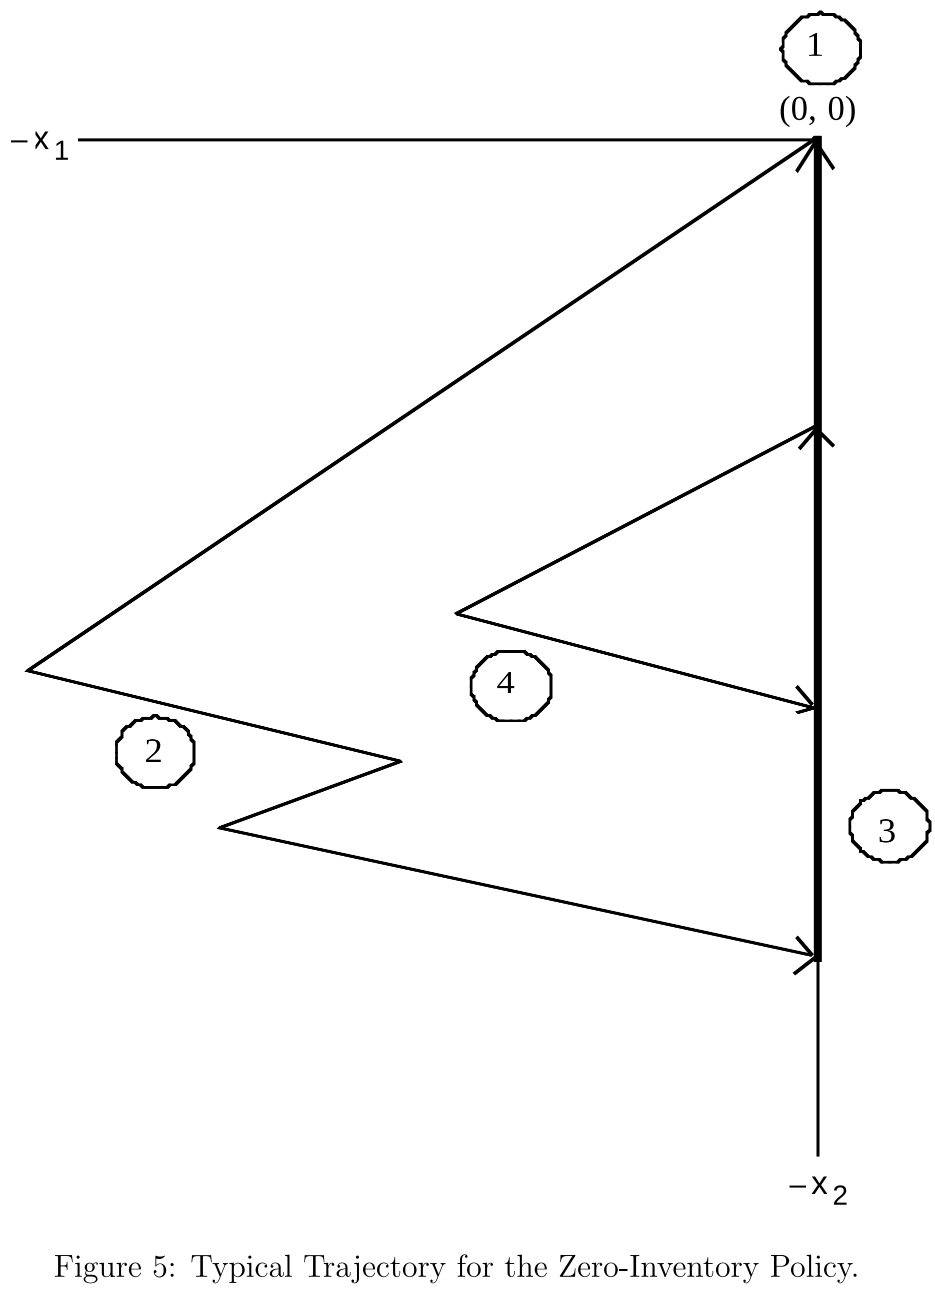

Figure 5: Typical Trajectory for the Zero-Inventory Policy.