Optimal Average Cost Manufacturing Flow Controllers: Convexity and Differentiability*

Michael H. Veatch,† *Gordon College*
Michael C. Caramanis,‡ *Boston University*

April 15, 1998

Abstract

We consider the control of a production facility consisting of a single workstation with multiple failure modes and part types using a continuous flow control model. Technical issues concerning the convexity and differentiability of the differential cost function are investigated. It is proven that under an optimal control policy the differential cost is C^1 on attractive control switching boundaries.

Index Terms Average cost minimization, value function differentiability, manufacturing flow control.

*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, mcaraman@bu.edu

1 Introduction

Manufacturing systems subject to discrete disturbances (failures, set-up changes and the like) have been studied extensively using a *fluid model* approximation, where surplus or backlog of production is represented by a continuous variable (see [5] for justification). The goal is to control production with a state feedback policy that minimizes the average cost of production surplus and backlog under a constant demand rate and stochastic production capacity. Little is known about the structure of the *optimal* policy for systems involving more than one part type; see Srivatsan and Dallery [12] Perkins and Srikant [8] and Veatch and Caramanis [14] for some recent exceptions. Instead, algorithms have been developed to compute a reasonable control policy using infinitesimal perturbation analysis or direct computation of average cost [2, 6, 7]. However, some of these algorithms rely on properties of the differential cost functions that have not been rigorously proven. Sethi, et al. [10] prove the existence of the potential cost function that is closely related to the differential cost.

This note investigates the continuity of the differential cost function's derivative on control switching surfaces, which are hyper-surfaces in the state space that form the boundaries between state space regions characterized by a constant optimal control. We show that the differential cost is, at least in some cases, continuously differentiable, justifying the assumption made in some previous papers and supporting the quadratic approximation used in [2]. Convexity of the differential cost is also established.

2 The Flow Control Model

We consider the flow control model of Liberopoulos and Caramanis [6], which generalizes the multiple unreliable machine model of [2]. The system state is $(x(t), \alpha(t))$, where $x = (x_1, \ldots, x_n)$, 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

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

To simplify notation, production constraints will be stated in terms of velocities $v(t) = \dot{x}(t)$. Let $\mathcal{V}_\alpha$ be the set of feasible velocities $v = \dot{x}$ in state α. For example, a single machine with maximum production rate $\overline{u}_i$ for type i has production constraint $\sum_i u_i/\overline{u}_i \le 1$. The feasible velocities are $\mathcal{V}_1 = \{v : \sum_i (v_i + d_i)/\overline{u}_i \le 1,\ v_i \ge -d_i\}$ in the "up" state and $\mathcal{V}_0 = \{-d\}$ in the "failed" state . We assume that the $\mathcal{V}_\alpha$ are convex polyhedra that satisfy the following:

A1. *Feasibility:* There exist $v^\alpha \in \mathcal{V}_\alpha$ such that $\sum_{\alpha \in \mathcal{E}} v^\alpha \pi_\alpha > d$

A2. *Probability Mass:* There exists at least one machine state, β^f which is non-transient, (i.e., $\pi_{\beta^f} > 0$) and satisfies $0 \in \mathcal{V}_{\beta^f}$.

(A1) guarantees that the problem is feasible, and it is essential for the stability of the controlled process. (A2) guarantees that under the optimal policy there exists a state $s = ((z_1, \ldots, z_n), \beta^f)$ at which the optimally controlled stochastic process has a probability mass [6, 8, 10]. The point $z^{\beta^f} = (z_1, \ldots, z_n)$ is in the production surplus space and is called the hedging point of machine state β^f.

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

$$g(x) \le C\left(1 + \sum_{i=1}^{n} |x_i|^\kappa\right) \tag{2}$$

for all x. The objective is to minimize long-run average cost. A control policy is the process $\{v(t) : t \ge 0\}$. 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 $t \ge 0$, 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 $v(t) = v(x(t), \alpha(t))$ ia a mapping of (x, α) to $\mathcal{V}_\alpha$ and does not depend on time explicitly. However, because $\alpha(\cdot)$ is memoryless and we are considering long-run average cost, methods such as [13] can be used to show that it is equivalent to optimize over all admissible policies. The control problem is

$$\min_{v} \quad \limsup_{T\to\infty} \frac{1}{T} E_{x,\alpha} \int_0^T g(x(t))dt \tag{3}$$

$$\text{s.t.} \quad \dot{x}(t) = v(t) \tag{4}$$

$$v \in V_M, \tag{5}$$

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

$$J^v(x, \alpha; T) = E_{x,\alpha} \int_0^T g(x(t))dt, \tag{6}$$

the long-run average cost of policy v by

$$\overline{J}^v = \limsup_{T\to\infty} J^v(x, \alpha; T)/T, \tag{7}$$

the optimal long-run average cost in (3-5) by $\overline{J}^*$, the differential cost for policy v by

$$W^v(x, \alpha) = \lim_{T\to\infty} J^v(x, \alpha; T) - T\overline{J}^v, \tag{8}$$

and the differential cost for the optimal policy by $W(x, \alpha)$ when these limits exist. The limits (7) and(8) may not exist for all policies.

The usual formulation of the HJB equations for this problem [6] assumes that W is continuously differentiable everywhere, which has not been shown. To avoid this difficulty, we make the weaker assumption that W has *one-sided directional derivatives*, justified in Section 4. Adopt the convention that

$$D_v f(x) = \lim_{h\downarrow 0} \frac{f(x + hv) - f(x)}{h}, \tag{9}$$

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

$$\overline{J}^* = g(x) + \min_{v\in\mathcal{V}_\alpha} D_v W(x, \alpha) + \sum_{\beta\in\mathcal{E}} q_{\alpha\beta} W(x, \beta) \tag{10}$$

for all $\alpha \in \mathcal{E}$. An informal derivation of (10), such as [2] or [4, Section 8.8], can be adapted to assume only that W has one-sided derivatives. We will also use the dynamic proramming equation

$$\overline{J}^v = g(x) + D_{v(x,\alpha)} W^v(x,\alpha) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} W^v(x,\beta), \tag{11}$$

which is just (10) without the minimization and is valid for any policy $v \in V_M$ for which $\overline{J}^v$, W^v, and $D_{v(x,\alpha)} W^v(x,\alpha)$ exist.

We derive another dynamic programming equation for later use. Let $x^v(t)$ denote the trajectory under policy v assuming no transitions occur. For small values of t, (6) for $J^v(x,\alpha;T)$ can be expanded to first order in t as,

$$J^v(x,\alpha;T) = g(x)t + J^v(x^v(t),\alpha;T-t) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} t J^v(x,\beta;T-t) + o(t).$$

Subtracting $T\overline{J}^v$ from both sides of the equation, taking the limit as $T \to \infty$, and using the fact that $\sum_{\beta \in \mathcal{E}} q_{\alpha\beta} = 0$, we obtain

$$W^v(x,\alpha) = \left[g(x) - \overline{J}^v \right] t + W^v(x^v(t),\alpha) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} t W^v(x,\beta) + o(t). \tag{12}$$

Now, $x^v(t)$ is Lipschitz continuous, and g is convex and, therefore, Lipschitz continuous; hence, the first term in (12) can be written $g(x)t = g(x^v(t))t + o(t)$. If W^v is Lipschitz continuous, the same argument applies to the last term in (12) and

$$W^v(x,\alpha) = \left[g(x^v(t)) - \overline{J}^v \right] t + W^v(x^v(t),\alpha) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} t W^v(x^v(t),\beta) + o(t). \tag{13}$$

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.

A3. *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 $\mathcal{V}_\alpha$.

A4. *Failed state:* For some $\alpha \in \mathcal{E}$, say $\alpha = 0$, $\mathcal{V}_0 = \{-d\}$, where $d > 0$.

3 Convexity

Convexity of the optimal differential cost W is established by noting that a) W exists since it is a solution to (10) which has been shown to exist [10], b) W differs from the potential or relative value function – defined below – by a constant, and c) the potential function is convex [10]. Convexity is used in the next section to prove that W has one-sided directional derivatives. The results in [10] and points a), b), and c) are summarized for the readers' convenience in the following theorem.

Theorem 1 *$W(x, \alpha)$ is convex in x.*

Proof. We need to show that $W(x, \alpha)$ differs by a constant from the potential cost function which Sethi, et al. [10] have proven to be convex.

Defining the cost-to-go function of the corresponding infinite horizon discounted cost problem for polict v as $J_\rho^v(x, \alpha) = \lim_{T\to\infty} E_{x,\alpha} \int_0^T e^{-\rho t} g(x(t)dt$ and optimal cost to go function as $J_\rho(x, \alpha) = \inf_{v\in V_M} J_\rho^v(x, \alpha)$, the average cost $\bar{J}$ and the potential function $V(x, \alpha)$, can be obtained as the following limits [10]:

$$\bar{J}^* = \lim_{\rho\to 0} \rho J_\rho(x, \alpha); \quad \bar{J}^v = \lim_{\rho\to 0} \rho J_\rho^v(x, \alpha)$$

$$V(x, \alpha) = \lim_{T\to\infty} [J_\rho(x, \alpha) - J_\rho(z^{\beta^f}, \beta^f)]; \quad V^v(x, \alpha) = \lim_{T\to\infty} [J_\rho^v(x, \alpha) - J_\rho^v(z^v, \beta^f)]$$

Sethi, et al. [10] prove that $\bar{J}^*$ and $V(x, \alpha)$ exist and are a solution to (10), and hence if v is the optimal policy, $\bar{J}^v$ and $V^v(x, \alpha)$ also exist and are a solution to (11).

Note that since $\sum_{\beta\in\mathcal{E}} q_{\alpha\beta} = 0$ if $V(x, \alpha)$ is a solution to (13), then $V(x, \alpha) + C$ is also a solution to (13) for any arbitrarily selected value of the constant C . After rearrangement, (13) becomes

$$\overline{J}^v = g(x + vt) + [W^v(x + vt, \alpha) - W^v(x, \alpha)] /t + \sum_{\beta\in\mathcal{E}} q_{\alpha\beta} W^v(x^v(t), \beta) + o(t). \tag{14}$$

Taking the limit as $t \to 0$, (14) implies (11) for a stable policy v and (10) for the optimal policy. Since both $V(x, \alpha)$ and $W(x, \alpha)$ must satisfy (10) or equivalently (13), it follows that

$V(x,\alpha)$ and $W(x,\alpha)$ differ by a constant. Hence, since $V(x,\alpha)$ is convex [10], it follows that $W(x,\alpha)$ is also convex.

□

Convexity of the cost function and theorem 1 imply that for any sample path of $\alpha(t)$ and for any initial states x^0, x^1 and $x^c = cx^0 + (1-c)x^1$, where $c \in (0,1)$, we have:

$$\int_0^T g(x^c(t))dt \leq c\int_0^T g(x^0(t))dt + (1-c)\int_0^T g(x^1(t))dt. \tag{15}$$

$$W(x^c,\alpha) \leq cW(x^0,\alpha) + (1-c)W(x^1,\alpha) \tag{16}$$

Strict convexity does not always hold. For example, in a single-machine problem with linear shortage costs, $g(x) = \sum_i g_i^- x_i^-$ for $x \leq 0$, and part types 1 and 2 tied under a "$c\mu$" rule, $g_1^-\overline{u}_1 = g_2^-\overline{u}_2$. Define the *workload* of types 1 and 2 as $w = x_1/\overline{u}_1 + x_2/\overline{u}_2$. Consider the coupled processes $x^0(t)$ and $x^1(t)$ with initial $x_i^0 = x_i^1, i \neq 1,2$, $w^0 = w^1$ and $x_i^j < 0, i = 1,2$. Use the optimal policy for $x^0(t)$. Also use this policy to set the type $i \neq 1,2$ production rates for $x^1(t)$, then use the optimal allocation between types 1 and 2. Call this policy v^1. Both policies have the same production capacity available for types 1 and 2 and will use all of the available capacity, maintaining equal workloads $w(t)$, as long as there are type 1 or 2 shortages. Furthermore, these policies eliminate backorders of types 1 and 2 before producing a surplus of either. Hence, the processes will merge before either has $x_1 > 0$ or $x_2 > 0$. We can write types 1 and 2 shortage cost as $g_1^- x_1 + g_2^- x_2 = g_1^-\overline{u}_1 w$, a function of w, showing that $g(x^0(t)) = g(x^1(t))$ and $J^*(x^0,\alpha;T) = J^{v^1}(x^1,\alpha;T) \geq J^*(x^1,\alpha;T)$. A symmetrical argument shows that $J^*(x^0,\alpha;T) \leq J^*(x^1,\alpha;T)$, demonstrating that v^1 must also be optimal. We have shown that $W(x^0,\alpha) = W(x^1,\alpha)$ on a line segment of equal workload with $x_1 \leq 0$ and $x_2 \leq 0$. Hence, equality (linearity) holds in (16) on this segment.

This example shows that strict convexity does not hold in the shortage region. Strict convexity applies if the cost rate g is strictly convex. The following corollary shows that it usually holds in the surplus region.

Corollary 1 *Suppose there is a failed state. Then $W(x,\alpha)$ is strictly convex on $\mathbf{R}_+^n$.*

Proof. Consider initial states $x^0 \neq x^1$ with $x^j \geq 0$, ordered so that $x_i^0 < x_i^1$ for some i. There is a nonzero probability of transitioning to the failed state $\alpha = 0$ while $x_i^0(t) < x_i^1(t)$ and, since $x^1(t) - x^0(t)$ is constant while in the failed state, there is also a nonzero probability of remaining in the failed state until $x_i^0(t) < x_i^1(t) = 0$. For such a sample path of $\alpha(t)$, strict inequality holds in (15) because g is nonlinear at $x_i = 0$. Then strict inequality holds for W on the line segment from x^0 to x^1 as well. □

4 Differentiability

It is frequently assumed that the differential cost W is continuously differentiable. This property simplifies notation, allows stronger characterizations of the optimal policy, and helps justify quadratic approximation of W. Rishel [9] shows, for a related class of problems with terminal rewards, that W^v is continuous everywhere and continuously differentiable where the control $v(x, \alpha)$ is continuous. Sethi, et al. [10] use the vanishing discount approach to show that the same result holds for the infinite horizon average cost problem considered here. Thus, the question of differentiability arises only on the *control switching sets* (CSS) where the control is discontinuous (see [2], [6], and [14]). We conjecture that W^v is continuously differentiable for a broad class of policies; however, we are only able to establish differentiability for the optimal policy and in certain situations. We also prove a weaker condition, that W has one-sided directional derivatives.

In all cases we are aware of where W^v is known, it is continuously differentiable. For the single part type, single-machine problem of [1], Gershwin ([4] p.286) gives the following argument for differentiability of W^v for any hedging point policy v. The control is only discontinuous at the hedging point z and $\alpha = 1$ (the machine is up), so we need only check $W^v(z, 1)$. The key observation is that, since all other terms in (11) are continuous, $D_{v(x(t),\alpha)} W^v(x(t), \alpha) = dW^v(x(t), \alpha)/dt$ is continuous in x. Because it is also a time derivative, it is zero at $(x(t), \alpha) = (z, 1)$ where $v(z, 1) = 0$. Approaching the hedging point from both sides and suppressing the policy superscript,

$$\lim_{x \to z^-} W'(x, 1)(\overline{u} - d) = \lim_{x \to z^+} W'(x, 1)(-d) = 0,$$

and $W'(z,1) = 0$. Our check of the formulas for W in [1] confirmed this result. The condition $W'(z,1) = 0$ for a suboptimal hedging point implies that $W(x,1)$ is not convex (we found this to be true in Gershwin's formulas for suboptimal $z = 0$), since otherwise the optimality conditions (10) would be satisfied.

In the two-part type problem with a single reliable machine, the optimal W found by Connelly [3] is continuously differentiable (after correcting algebraic errors). For the unreliable version of this problem, Srivatsan [11] uses the optimal policy v for the deterministic system and suggests that W^v is not differentiable on the control region boundaries; however, we suspect that this is an erroneous conclusion due to errors in [3].

Returning to the general problem, we begin with the one-sided directional derivative defined in (9). The following results are stated for the optimal policy so that convexity of W can be used; however, we suspect that a suitably restricted class of policies would imply properties of W^v that lead to the same results.

Lemma 1 *W has one-sided directional derivatives.*

Proof. For a given x, v and $\delta > 0$, the difference quotient

$$\frac{W(x + hv, \alpha) - W(x, \alpha)}{h}$$

is bounded above and below for all $h < \delta$ because W is convex. It decreases as h decreases for the same reason. Hence, the limit as $h \downarrow 0$ exists. □

In light of Rishel's results, the next question is whether W is differentiable at boundaries where constant control regions intersect. These boundaries are Control Switching Surfaces (CSS) which can be formally defined as follows: For a given machine state α, consider a set of allowable controls $\mathcal{V}_\alpha{}^* \subseteq \mathcal{V}_\alpha$. There is a CSS which corresponds to a given $\mathcal{V}_\alpha{}^*$. That CSS is the subset of the production surplus space X whose elements x achieve the minimum in (3-5) for any control $v \in \mathcal{V}_\alpha{}^*$. Depending on the order of singularity characterizing the elements of X, $\mathcal{V}_\alpha{}^*$ may contain a single control (i.e., the optimal policy is unique and the order of singularity is zero) or infinite controls (i.e., the optimality conditions (10) are singular and satisfied by an uncountable number of controls). In any case, $\mathcal{V}_\alpha{}^*$ can be expressed as the convex hull of a minimal set of extreme points of $\mathcal{V}_\alpha$. As elaborated below, the dimension of

${\mathcal{V}_\alpha}^*$ increases as the dimension of X decreases. For example, for ${\mathcal{V}_\alpha}^* = \mathcal{V}_\alpha$, the corresponding X may be countable or even empty (except when $\alpha = 0$ and $\mathcal{V}_\alpha = \{-d\}$. Let Conv denote the convex hull and define $\{v^1, \ldots, v^l\}$ to be the minimal set of extreme points of $\mathcal{V}_\alpha$, such that ${\mathcal{V}_\alpha}^* \subseteq \text{Conv}(v^1, \ldots, v^l)$. The CSS that corresponds to ${\mathcal{V}_\alpha}^*$ can be considered as a function of $\{v^1, \ldots, v^l\}$ and denoted as $X(v^1, \ldots, v^l)$, where

$$X(v^1, \ldots, v^l) = \{x : \{v^1, \ldots, v^l\} \text{ is minimal s.t. } \mathcal{V}_\alpha^*(x) \subseteq \text{Conv}(v^1, \ldots, v^l)\}. \tag{17}$$

We next proceed to describe further the CSSs. For the single-machine problem of [6], 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. 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 $0 \notin \mathcal{V}_\alpha$, i.e., in a machine state that is not able to meet demand, $x(t)$ 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. For an α with no hedging point, the trajectory may also cross a boundary that removes more extreme points from the optimal set than are added. Assumption (A2) guarantees that there is at least one state α^f with a hedging point.

We make two adjustments to the structure in [6]: (i) because we do not assume that W is continuously differentiable, (10) may not have extreme point solutions for all x and (ii) a system can be stable without having a hedging point (for example when assumption (A1) holds but (A2) does not). If W is not differentiable at x, then $\mathcal{V}_\alpha^*(x)$ may not contain an extreme point of $\mathcal{V}_\alpha$; instead, it may be in the interior of some face $\text{Conv}(v^1, \ldots, v^l)$ of $\mathcal{V}_\alpha$.

An attractive CSS, in the sense of [6], has the property that it is possible to remain in the CSS for a nonzero amount of time using a control in $\text{Conv}(v^1, \ldots, v^l)$. Roughly, an attractive CSS extends in a direction that is in the convex cone of the vectors $v^1, \ldots, v^l$. A CSS is also attractive if 0 is a convex combination of the vectors $v^1, \ldots, v^l$. 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 [6]. 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.

Theorem 2 *If $x \in X(v^1, v^2)$, has the attractive property and $X(v^1, v^2)$ is a surface of dimension $n-1$ in $\mathbf{R}^n$, then $W(\cdot, \alpha)$ is differentiable at x.*

Proof. By the definition of attractive CSSs, there is a control

$$b = cv^1 + (1-c)v^2 \tag{18}$$

that lies in the direction of $X(v^1, v^2)$, in the sense that the distance of the point $x + tb$ from the CSS $X(v^1, v^2)$ is $o(t)$, for otherwise the trajectory $x(t)$ could not remain in $X(v^1, v^2)$. If $X(v^1, v^2)$ possesses a tangent hyperplane at x, b lies in the hyperplane. Since W is convex and continuously differentiable on the CSS $X(v^i)$, $i = 1, 2$, $\nabla W(x_i, \alpha)$ has a limit $\nabla_i W$ as $x_i \to x$ through any sequence of points in $X(v^i)$. Because $X(v^1, v^2)$ is of lower dimension, x is a limit point of $X(v^i)$. We will show that $\nabla_1 W = \nabla_2 W = \nabla W(x, \alpha)$.

ADD FIGURE ATTACHED AT THE END OF THE PAPER HERE.

For directions v from x that lie in $X(v^i)$, $D_v W(x, \alpha) = \nabla_i W \cdot v$. The continuity of W requires that the two derivatives match in the boundary direction:

$$\nabla_1 W \cdot b = \nabla_2 W \cdot b = D_b W(x, \alpha). \tag{19}$$

Since W is continuous, the other terms in (10) must be also, which requires

$$\nabla_1 W \cdot v^1 = \nabla_2 W \cdot v^2 = D_b W(x, \alpha). \tag{20}$$

Combining (18)-(20),

$$\begin{aligned}
\nabla_1 W \cdot v^2 &= \nabla_1 W \cdot \left[v^1 + \frac{1}{1-c}(b - v^1) \right] \\
&= \nabla_1 W \cdot v^1 \\
&= \nabla_2 W \cdot v^2.
\end{aligned}$$

Now we can construct a basis for $\mathbf{R}^n$ consisting of v^2 and $b^1, \ldots, b^{n-1}$, where each b^i is in the direction of the surface $X(v^1, v^2)$ at x. Then (19) holds with b^i in place of b, and equality of the dot products on a basis implies $\nabla_1 W = \nabla_2 W$. □

Next we establish differentiability of W at the hedging point. The argument is an extension of Gershwin's argument for the single part type problem.

Theorem 3 *For machine states α with 0 an interior point of $\mathcal{V}_\alpha$, $W(\cdot, \alpha)$ is differentiable at the hedging point z^α.*

Proof. Consider a trajectory using the optimal policy beginning at (x, α). If there are no changes of state, the trajectory will reach z^α at some time t. At $x \neq z^\alpha$ the optimal control v is a boundary point of $\mathcal{V}_\alpha$, so $\|v\| \geq m > 0$ for some m that does not depend on x. Also, the length of the trajectory from x to z^α is less than $K\|x - z^\alpha\|$ for some K that depends only on the geometry of $\mathcal{V}_\alpha$. Hence, $t \leq K\|x - z^\alpha\|/m$. Using (13) at this t, we obtain

$$W(x, \alpha) = [g(z^\alpha) - J]\, t + W(z^\alpha, \alpha) + \sum_{\beta \in \mathcal{E}} q_{\alpha\beta} t W(z^\alpha, \beta) + o(t).$$

But a trajectory starting at z^α also has $x(t) = z^\alpha$ if there are no transitions, so the same expression holds for $W(z^\alpha, \alpha)$, and $W(x, \alpha) - W(z^\alpha, \alpha) = o(t) = o(\|x - z^\alpha\|)$. Letting $x \to z^\alpha$ from the direction v, $D_v W(z^\alpha, \alpha) = 0$. Since this holds for all v, $\nabla W(z^\alpha, \alpha) = 0$. □

For the two part type single-machine problem, these theorems cover all cases and W is continuously differentiable. It seems as if it should be possible to extend Theorem 2 to boundaries of more than two control regions.

Acknowledgements

We thank Stan Gershwin and Jim Perkins for their helpful comments and Inna Zharnitsky for her assistance in checking differentiability.

References

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

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

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

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

[5] 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.

[6] 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.

[7] J. R. Perkins, and R. Srikant, "Scheduling Multiple Part-Types in an Unreliable Single Machine Manufacturing System," *IEEE Trans. Automat. Contr.,* Vol. 42, No 3, pp. 364-377, 1997.

[8] J. R. Perkins, and R. Srikant, "Hedging Policies for Failure- Prone Manufacturing Systems: Optimality of JIT and Bounds on Buffer Levels," *IEEE Trans. Automat. Contr.,* Vol. 43, 1998, in press.

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

[10] 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 Dynamic Systems: Theory and Applications,* Vol 8, 1998, in press.

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

[12] N. Srivatsan and Y. Dallery, "Partial Characterizations of Optimal Hedging Point Policies in Unreliable Two-Part-Type Manufacturing Systems," *Oper. Res.*, Vol. 46, No. 1, 1998, in press.

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

[14] M. H. Veatch and M. Caramanis, "Optimal Manufacturing Flow Controllers: Zero-Inventory Policies and Control Switching Sets", Boston University Dept of Manuf Engineering Working paper, 1996, Submitted to *IEEE Trans. Automat. Contr.*.