Managing Capacity by Drift Control

Melda Ormeci Matoglu[1], John Vande Vate[2],

[1] Ozyegin University, Faculty of Economics and Administrative Sciences
Kusbakisi Cad. No 2, 34662, Altunizade, Istanbul, TURKEY
melda.ormeci@ozyegin.edu.tr
[2] Georgia Institute of Technology, School of Industrial and Systems Engineering
Atlanta, Georgia, 30332-0205, USA
john.vandevate@isye.gatech.edu

Abstract. We model the problem of managing capacity in a build-to-order environment as a Brownian drift control problem and seek a policy that minimizes the long-term average cost. We assume the controller can, at some cost, shift the processing rate among a finite set of alternatives. The controller incurs a cost for capacity per unit time and a delay cost that reflects the opportunity cost of revenue waiting to be recognized or the customer service impacts of delaying delivery of orders. Furthermore he incurs a cost per unit to reject orders or idle resources as necessary to keep the workload of waiting orders within a prescribed range. We introduce a practical restriction on this problem, called the $\mathcal{S}$-restricted Brownian control problem, and model it via a structured linear program. We demonstrate that an optimal solution to the $\mathcal{S}$-restricted problem can be found among a special class of policies called deterministic non-overlapping control band policies. These results rely on (apparently new) relationships between complementary dual solutions and relative value functions, through which we to obtain a lower bound on the average cost of any non-anticipating policy for the problem even without the $\mathcal{S}$ restriction. Under mild assumptions on the cost parameters, we show that our linear programming approach is asymptotically optimal for the unrestricted Brownian control problem in the sense that by appropriately selecting the $\mathcal{S}$-restricted problem, we can ensure its solution is within an arbitrary finite tolerance of a lower bound on the average cost of any non-anticipating policy for the unrestricted Brownian control problem.

Keywords. Capacity Management, Brownian Motion, LP, Relative Value Function, Duality

1 Introduction

We consider the problem of managing capacity in a build-to-order environment like the personal computer manufacturing operations of Dell. We model this problem as a Brownian control problem and seek a policy that minimizes the long-run average cost. The workload of orders awaiting processing incurs a delay cost or waiting cost per unit per unit time that reflects the opportunity cost

Dagstuhl Seminar Proceedings 09261
Models and Algorithms for Optimization in Logistics
http://drops.dagstuhl.de/opus/volltexte/2009/2159

of revenue waiting to be recognized or the customer service impacts of delaying delivery of orders. On the other hand, maintaining capacity at a higher production rate also incurs a higher cost for capacity per unit time. We assume the controller can, at some cost, shift the processing rate among a finite set of alternatives by, for example, adding or removing staff, increasing or reducing the number of shifts or opening or closing production lines. We model the cost of changing the processing rate from u to v as a fixed cost $K(u,v) > 0$ and assume the cost function K satisfies the usual triangle inequality so that changing the processing rate from u to v in a single step is no more expensive than accomplishing the same change via a series of intermediate steps.

Even with changes in the processing rate, the workload of waiting orders can grow without limit. To ensure that delivery commitments can be met, we introduce additional *boundary controls*. In particular, we impose a fixed maximum level on the workload of waiting orders and exert instantaneous controls, e.g., rejecting orders, to keep from exceeding this limit.

Similarly, the workload of waiting orders can drain to 0 and to ensure that the process remains non-negative, i.e., that the manufacturer does not build product in advance of orders, we impose instantaneous controls at the lower boundary corresponding roughly to idling lines briefly until additional orders arrive.

In this paper, we consider a restricted version of the problem, called the *$\mathcal{S}$-restricted Brownian control problem*, in which the controller may change the drift rate only when the workload of waiting orders reaches a value in a given finite set $\mathcal{S}$. Choosing $\mathcal{S}$ to be the set of non-negative integers up to the maximum workload for example, imposes no limitation on the system in practice, but allows us to formulate a linear program approximating the problem by considering only a special class of policies called control band policies. We show that an optimal solution to our structured linear program can be found among the special class of deterministic non-overlapping control band policies and that an optimal solution to the linear program actually solves the $\mathcal{S}$-restricted Brownian control problem, i.e., a deterministic non-overlapping control band policy is optimal among all non-anticipating policies for the $\mathcal{S}$-restricted problem. We further provide a lower bound on the average cost of any non-anticipating policy (i.e., any policy whose decisions at time t depend only on the history up to time t, and not on future events) for the unrestricted problem and show that under mild assumptions on the cost parameters the average cost of an optimal solution to the restricted problem can be made arbitrarily close to this lower bound.

Drift control problems were studied in the literature in different contexts with different cost structures and solution approaches. See, for example, [1,2,3,4,5,6,7]. A major difference between the works of [2,3,6,7] and ours is that these works restrict the controller to only two drift rates. In [1], the authors solve a drift control problem in which the process is confined to a finite range by instantaneous controls at the boundaries and the objective is to minimize the long term average cost of control for drift and displacement at the upper boundary. They show that the optimal drift rate in each state is the smallest minimizer of the Bellman equation they derive. In [4] the authors address the same problem, with the added task of determining the optimal buffer size. A major difference between these works and ours is that these works restrict the controller to only two drift rates. The model of [1] includes more general processing costs, but does not address the holding or delay costs and changeover costs, in [4] they include a congestion cost similar to a holding cost but do not include changeover costs. The changeover costs in our model make the controller liable for past decisions and result in optimal policies that depend not only on the position of the process but also on the current drift rate.

In [7] and [3] the authors study a reflected Brownian motion process in which the controller can switch between two sets of drift and diffusion parameters. The problem involves operating, switching and holding costs. Since only a reflective boundary at zero was defined, at least one of the drift rates must be negative. While in [7] the authors solve the continuous time control problem by approximating the Brownian motion processes using the corresponding random walks, in [3] the authors address the same problem directly by treating it in continuous time using dynamic programming. The optimal policy is similar to our control band policies, with two switching points. In [6] they address a similar problem where a controller chooses between two drift rates, however they take the form of the policy as given and look at several different cost functions.

Our model differs from the ones in these works in many ways. We address the more general problem of selecting from many rates. Furthermore, instead of a reflecting boundary at zero with zero cost, we require the process to remain between two fixed boundaries and impose costs on the instantaneous controls needed to keep it there. The process incurs the cost of lost production whenever the lines are idled and the cost of rejecting orders whenever the upper boundary is reached, reflecting both the immediate lost revenue and the potential impact on future sales to the customer.

Our work also differs from the ones in the literature as we develop a novel solution approach based on linear programming and exploiting relationships between dual solutions and relative value functions.

Our linear programming approach directly addresses the continuous time Brownian control problem by isolating the individual cost components and explicitly calculating the frequencies at which controls are exerted. Being able to quickly obtain a breakdown of costs, and see the impacts of policy changes is a valuable tool to the controller, which the other approaches do not provide.

The use of linear programming to reformulate long-term average stochastic control problems began with [8] in the context of a discrete time, finite state controlled Markov chain and now has become standard (see, for example, [9]). In [10], general discretization schemes based on approximate controlled Markov chains are introduced to solve stochastic control problems in continuous time and continuous state. In [11], a discretization scheme using the finite element method is developed for certain singular control diffusion problems. This method performs better than those in [10], but again does not provide an error bound with respect to the continuous problem. In [12] and [13], linear programming based approaches are developed to solve diffusion control problems. These methods generate constraints on a finite set of moments to develop an approximate solution, but do not provide an error bound.

In this paper we develop an innovative discretization scheme based on drift rates and transition points that not only provides a near-optimal average cost, but also explicitly defines an easily implementable policy to achieve that average cost. The discretization scheme is natural to and consistent with industrial settings. We provide error bounds on the quality of the solutions produced via our discretization scheme and show that these solutions can be made asymptotically optimal by appropriate refinements of the discretization.

In Section 2, we describe the average cost Brownian control problem and its policy space. In Section 3 we state our linear programming formulation and main results.

2 Brownian Control Problem

Let

$$W(T) = W(0) + \int_0^T \mu(t)dt + \sigma B(T), \quad T \geq 0, \tag{1}$$

be a diffusion process with drift $\mu(t)$ in some fixed finite set Λ for each $t \geq 0$, variance $\sigma^2 > 0$ and initial level $W(0)$ on some filtered space $\{\Omega, \mathcal{F}, \mathbb{P}; \mathcal{F}_t, t \geq 0\}$. The process $W(T)$ describes the difference between cumulative work to have arrived and cumulative work processed by time T, i.e. the netput process. The drift rate $\{\mu(t), t \geq 0\}$, which is adapted to the Brownian motion $\{B(t) : t \geq 0\}$, is the difference between the average arrival rate and the rate $\lambda(t)$ at which work is completed. We assume the arrival process is time homogeneous with average rate μ_0 and that the controller can, at some cost, shift the processing rate among a finite set of alternatives. Further, the controller must exert the minimal instantaneous control required to keep the process within the allowed range $[0, \Theta]$. We let $A(t)$ denote the cumulative increases in work and $R(t)$ the cumulative decreases in work up to time t exerted by the controller at 0 and Θ, respectively. The resulting controlled process is

$$X(T) = X(0) + \int_0^T \mu(t)dt + \sigma B(T) + A(T) - R(T), \quad T \geq 0, \tag{2}$$

where $X(0) = W(0)$. The controlled process $X(t)$ lives in the bounded region $[0, \Theta]$, and the controller may only adjust the drift rate by choosing from among the possible values in the finite set Λ. We assume, without loss of generality, that $W(0) \in [0, \Theta]$. To avoid tedious case analysis, we also assume that $0 \notin \Lambda$. Analogous results hold when $0 \in \Lambda$. We let $\mathcal{D} = \{(a, u) : a \in [0, \Theta], u \in \Lambda\}$ denote the domain of this process.

A policy defines the times at which and amounts by which we adjust the drift rate. We restrict attention to the space $\mathcal{P}$ of all non-anticipating policies $\Phi = \{(T_i, u_i) : i \geq 0\}$, where (i) $0 \leq T_0 < T_1 < T_2 < ... < T_i < T_{i+1}, \ldots$ is a sequence of stopping times and (ii) Each $u_i \in \Lambda$ is a random variable adapted to $\mathcal{F}_{T_i}$ indicating the rate to which we change the drift at time T_i. Under the policy $\Phi = \{(T_i, u_i) : i \geq 0\}$, the drift rate $\mu(t) = u_i$ for $T_i \leq t < T_{i+1}$.

With each policy $\{(T_i, u_i) : i \geq 0\} = \{\mu(t) : t \geq 0\}$, the associated Skorohod problem:
(a) $X(t) \in [0, \Theta]$, $t \geq 0$,
(b) $A(\cdot)$, $R(\cdot)$ are nondecreasing and continuous with $A(0) = 0$, $R(0) = 0$,
(c) $\int_0^T 1_{\{X(t)>0\}} dA(t) = \int_0^T 1_{\{X(t)<\Theta\}} dR(t) = 0$, $t \geq 0$,
where the continuous process $\{X(t) : t \geq 0\}$ is defined by (2), uniquely defines $\{X, A, R\}$ (See Section 2.4 of [14]). Note that since the drift rate controls uniquely determine the instantaneous controls exerted at the boundaries, we do not include the latter in our specification of a policy.

To change the drift from rate u to rate v, the controller must pay a fixed cost, $K(u, v) > 0$, for $u \neq v$, which satisfies a triangle inequality:

$$K(u, v) + K(v, w) \geq K(u, w) \text{ for all rates } u, v \text{ and } w. \tag{3}$$

To simplify notation, we let $K(u, u) = 0$ for all $u \in \Lambda$.

There is a cost $c(u)$ per unit time for the capacity to process work that depends on the drift rate u and when $X(t) > 0$ there is a backlog of orders, which incurs a linear delay cost at rate h per unit per unit time. The instantaneous controls exerted at 0 and Θ to adjust the workload either up or down incur a unit cost of U and M, respectively.

We consider the *Average Cost Brownian Control Problem*, which is to find a non-anticipating policy that minimizes the long run average cost:

$$\mathrm{AC}(\Phi) = \limsup_{T\to\infty} \frac{1}{T} \mathbb{E}[\int_0^T c(\mu(t)) + hX(t)dt + UA(T) + MR(T) + \sum_{i=1}^{N(T)} K(u_{i-1}, u_i)] \tag{4}$$

where, $\mu(t)$ denotes the drift rate at time t and, for each $t \geq 0, N(T) = \sup\{n \geq 0 : T_n \leq T\}$ denotes the number of changes in the drift rate by time T.

A *control band* $\psi = (u, s, S)$ is defined by a rate $u \in \Lambda$ and an interval (s, S). Given $\mu(t) = u$ and $X(t) \in (s, S)$, a policy implementing the control band $\psi = (u, s, S)$ maintains the drift rate u until X first reaches $\{s, S\}$. Given a control band $\psi = (u, s, S)$, we let $u(\psi)$ denote the rate u, $s(\psi)$ denote the lower limit s and $S(\psi)$ denote the upper limit S. We say that a point $z \in [0, \Theta]$ is contained in ψ if $z \in (s, S)$. A point $(z, u) \in I\!R \times \Lambda$ is contained in the control band ψ if z is contained in ψ and $u = u(\psi)$.

The fact that $\mu(t)$ is right-continuous with left-hand limits leads to a certain ambiguity in the interpretation of a lower limit $s = 0$ or an upper limit $S = \Theta$. When $s = 0$, it is not clear whether we mean to change the drift rate or to invoke instantaneous control when $X(t)$ reaches 0. To resolve this ambiguity, we distinguish between a control band with lower limit $s = 0$, which changes the drift rate when $X(t) = 0$, and one with lower limit $s = -\infty$, which simply relies on instantaneous control at 0 to keep $X(t) \geq 0$ and maintains the drift rate until $X(t)$ first reaches S. Similarly, we distinguish between a control band with upper limit $S = \Theta$ and one with $S = \infty$, which simply relies on instantaneous control at Θ to keep $X(t) \leq \Theta$. Thus, the lower limit of a control band must be in the set $\{-\infty\} \cup [0, \Theta]$ and the upper limit must be in the set $[0, \Theta] \cup \{\infty\}$. More specifically, compare the control bands $\psi = (u, s, S)$ where: (i) $s = 0$ and the process switches to band ψ_1 if $X(t)$ reaches the lower limit s, (ii) $S = \Theta$ and the process switches to band ψ_2 if $X(t)$ reaches the upper limit S, (iii) $s = -\infty$ and (iv) $S = \infty$. Then

$$\mu(t+) = \begin{cases} u(\psi_1), & \text{if } X(t) = 0 \text{ and } s = 0, \\ u(\psi_2), & \text{if } X(t) = \Theta \text{ and } S = \Theta, \\ u, & \text{if } X(t) = 0 \text{ and } s = -\infty, \\ u, & \text{if } X(t) = \Theta \text{ and } S = \infty. \end{cases}$$

Given a collection $\Psi = \{\psi_i : i \in \mathcal{I}\}$ of control bands, we often refer to the band ψ_i simply by its index i, so for example, we refer to the band ψ_i simply as band i, to $u(\psi_i)$ as $u(i)$, etc. We say that two bands i and j *overlap* if $u(i) = u(j)$ and $(s(i), S(i)) \cap (s(j), S(j)) \neq \emptyset$. Otherwise, bands i and j are *non-overlapping*. We say that a collection of bands is non-overlapping if no two of its members overlap. Note that our definition of non-overlapping bands includes bands with the same drift rate whose intervals share a common endpoint. If a policy includes overlapping bands additional management tools may be required to resolve which band should be in effect at any given time.

A *control band policy* is defined by a collection of control bands $\Psi = \{\psi_i : i \in \mathcal{I}\}$ such that each point $x \in [0, \Theta]$ is contained in some band ψ_i together with a rule for switching from one control band to the next. A control band policy maintains the drift rate of the current control band ψ until $X(t)$ first reaches $s(\psi)$ or $S(\psi)$ at which point it changes to a new band as dictated by the switching rule and adopts the corresponding drift rate.

A *deterministic control band policy* is a control band policy defined by a collection $\Psi = \{\psi_i : i \in \mathcal{I}\}$ together with two maps $\zeta : \mathcal{I} \mapsto \mathcal{I}$ and $\tau : \mathcal{I} \mapsto \mathcal{I}$. The switching rule for this policy follows the map ζ when $X(t)$ reaches the (finite) lower limit of the current band and follows the map τ when $X(t)$ reaches the (finite) upper limit of the current band. In particular, if the policy is in control band i and $X(t)$ first reaches $s(i)$, the policy deterministically changes to the control band $\zeta(i)$ and so changes the drift rate to $u(\zeta(i))$. Similarly, if $X(t)$ first reaches $S(i)$, the policy deterministically changes to the control band $\tau(i)$ and so changes the drift rate to $u(\tau(i))$. Deterministic policies are simple to describe and implement.

If the process starts at a point $(X(0), \mu(0))$ that is not in any control band of the policy, we may transition to a band of the policy without affecting the average cost so long as this transition is accomplished in finite time and with finite cost. One method to accomplish this one-time transition is, for example, to maintain the drift rate $\mu(0)$ until the process first reaches a point in some specified set and then switch to any band of the policy that contains this point.

In this paper we consider the $\mathcal{S}$-restricted Brownian control problem in which the controller may only change the drift rate when $X(t)$ is in a given finite set $\mathcal{S} = \{s_i : i = 1, 2, \ldots, n\}$, where $0 < s_1 < s_2 \ldots < s_n < \Theta$. We define the subset $\mathcal{P}(\mathcal{S}) \subset \mathcal{P}$ to be those non-anticipating policies $\{(T_i, u_i) : i \geq 0\}$ in which each T_i is a hitting time for some subset of $\mathcal{S}$. Every control band policy in which the (finite) end points of each control band are in $\mathcal{S}$ is a non-anticipating policy of the $\mathcal{S}$-restricted Brownian control problem and so a member $\mathcal{P}(\mathcal{S})$. Not every policy in $\mathcal{P}(\mathcal{S})$, however, is a control band policy since, for example, the hitting times need not be first hitting times or may involve other complicating conditions.

3 The LP Formulation and Summary of Results

We formulate a linear program $\mathrm{LP}(\Lambda, \mathcal{S})$ to find a minimum average cost control band policy in $\mathcal{P}(\mathcal{S})$. The formulation of $\mathrm{LP}(\Lambda, \mathcal{S})$ requires the following performance characteristics of each control band as input.

Consider a band $\psi = (u, s, S)$ and a point $z \in (s, S)$. We let $p(z, \psi, s)$ denote the probability of exiting band ψ at s given the starting point z, i.e. the probability that, starting at z, $X(t)$ with drift rate fixed at u first reaches s. Thus, $p(z, \psi, S)$ is the probability of exiting band ψ at S given the starting point z. Observe that if both endpoints s and S of the band are finite, then $p(z, \psi, s) + p(z, \psi, S) = 1$. If only one endpoint of ψ is finite, the probability of exiting the band at that endpoint is clearly 1. If both endpoints are infinite, i.e., if $s = -\infty$ and $S = \infty$, then there are no transitions out of band ψ. Such bands are called *absorbing*.

We let $\mathbb{E}[X|(z, \psi)]$ denote the average value of $X(t)$ over the time the process is in band $\psi = (u, s, S)$ given we enter the band at point $z \in (s, S)$. When $s = -\infty$, we let $\mathbb{E}[A|(z, \psi)]$ denote the average rate per unit time at which the controller must exert instantaneous control to keep the process non-negative in band ψ when it enters the band at the point z. Similarly, when $S = \infty$ we let $\mathbb{E}[R|(z, \psi)]$ denote the average rate per unit time at which the controller must exert instantaneous control to keep the process from exceeding Θ in band ψ when it enters the band at the point z. Observe that $\mathbb{E}[A|(z, \psi)] = 0$ when $s(\psi) \geq 0$ and $\mathbb{E}[R|(z, \psi)] = 0$ when $S(\psi) \leq \Theta$.

We denote by $\mathrm{Cost}(z, \psi)$ the average rate at which the process accumulates cost in band $\psi = (u, s, S)$ when it enters the band at the point z. In particular,

$$\mathrm{Cost}(z, \psi) = c(u) + h\mathbb{E}[X|(z, \psi)] + U\mathbb{E}[A|(z, \psi)] + M\mathbb{E}[R|(z, \psi)].$$

The average rate at which the process accumulates cost in an absorbing band ψ with drift rate u is independent of the point at which it enters the band. In fact, the cost only depends on u. In this case, we often write $\mathrm{Cost}(u)$ in place of $\mathrm{Cost}(z, \psi)$.

We let $\mathbb{E}[T|(z, \psi)]$ denote the expected time for the process to hit $\{s, S\}$ when it enters band $\psi = (u, s, S)$ at the point z. We also use $\delta(z, \psi, s)$ to denote the average rate per unit time at which the process reaches s when it enters band ψ at the point z. Thus, $\delta(z, \psi, S)$ denotes the average rate per unit time at which the process reaches S when it enters band ψ at the point z. When $-\infty = s < S \leq \Theta$, $\delta(z, \psi, S)$ is defined to be $1/\mathbb{E}[T|(z, \psi)]$. Similarly, when

$0 \le s < S = \infty$, $\delta(z,\psi,s)$ is defined to be $1/\mathbb{E}[T|(z,\psi)]$. Finally, when both s and S are finite $\delta(z,\psi,s)$ and $\delta(z,\psi,S)$ are defined by the relations:

$$\begin{aligned}
1/\mathbb{E}[T|(z,\psi)] &= \delta(z,\psi,s) + \delta(z,\psi,S) \\
p(z,\psi,s) &= \frac{\delta(z,\psi,s)}{\delta(z,\psi,s)+\delta(z,\psi,S)} \\
p(z,\psi,S) &= \frac{\delta(z,\psi,S)}{\delta(z,\psi,s)+\delta(z,\psi,S)}.
\end{aligned}$$

These quantities may easily be obtained by solving a small system of linear equations derived through Basic Adjoint Relations (BAR) [15,16].

At this point we may formulate the linear program $\mathrm{LP}(\Lambda,\mathcal{S})$ approximating the $\mathcal{S}$-restricted Brownian control problem.

Given a finite set Λ of potential drift rates and a finite set $\mathcal{S}$ of candidate transition points we define $\Psi(\Lambda,\mathcal{S}) = \{\psi_i : i \in \mathcal{I}(\Lambda,\mathcal{S})\}$ to be the finite collection of control bands $\psi = (u,s,S)$ with $u \in \Lambda$ and s and S in $\mathcal{S} \cup \{-\infty,\infty\}$ such that $s < S$. The set $\mathcal{Z}(\Lambda,\mathcal{S}) = \{(z,i) : i \in \mathcal{I}(\Lambda,\mathcal{S}),\ z \in \mathcal{S} \cap (s(i),S(i))$ and ψ_i is not absorbing$\}$ indicates all the possible points $z \in \mathcal{S}$ at which the process can transition into each non-absorbing band i in $\Psi(\Lambda,\mathcal{S})$.

The linear programming formulation of the $\mathcal{S}$-restricted Brownian control problem, $\mathrm{LP}(\Lambda,\mathcal{S})$, has three types of variables:

Band Variables:

- $x(z,i)$ for each entry point $(z,i) \in \mathcal{Z}(\Lambda,\mathcal{S})$, describing the rate or average number of transitions per unit time at which the process enters the band i at the entry point z and,
- $x(u)$ for each $u \in \Lambda$, describing the long-run fraction of time that the absorbing band $(u,-\infty,\infty)$ is used. Since there are no transitions out of absorbing bands, $x(u)$ should be either 0 or 1, but it can be shown that it is not necessary to impose that restriction.

Transition Variables: $y(z,u,v)$ for each triple $(z,u,v) \in \mathcal{S}\times\Lambda\times\Lambda$, describing the rate per unit time at which the process changes the drift from u to v at point z.

w-Variables: $w(z,u)$ for each $(z,u) \in \mathcal{S}\times\Lambda$, which we include as a convenient mechanism for effectively allowing "$y(z,u,u)$" to be a free variable by representing it as the difference of two non-negative variables $w(z,u)$ and $y(z,u,u)$.

We denote the objective function value of a solution (x^*,y^*,w^*) by $\mathrm{AC}(x^*,y^*,w^*)$.

Minimize

$$\sum_{(z,i)\in\mathcal{Z}(\Lambda,\mathcal{S})} \mathrm{Cost}(z,i)\mathbb{E}[T|(z,i)]x(z,i) + \sum_{u\in\Lambda}\mathrm{Cost}(u)x(u) + \sum_{(z,u,v)\in\mathcal{S}\times\Lambda\times\Lambda} K(u,v)y(z,u,v)$$

subject to the constraints:

Scale Constraint:

$$\sum_{(z,i)\in\mathcal{Z}(\Lambda,\mathcal{S})} \mathbb{E}[T|(z,i)]x(z,i) + \sum_{u\in\Lambda} x(u) = 1 \quad (5)$$

In-Conservation Constraint for each $(z,v) \in \mathcal{S}\times\Lambda$:

$$\sum_{(z,i)\in\mathcal{Z}(\Lambda,\mathcal{S}):\, u(i)=v} x(z,i) + w(z,v) - \sum_{u\in\Lambda} y(z,u,v) = 0 \quad (6)$$

Out-Conservation Constraint for each $(z,u) \in \mathcal{S}\times\Lambda$:

$$-\sum_{\substack{(z',i)\in\mathcal{Z}(\Lambda,\mathcal{S}):\\ u(i)=u \text{ and } z\in\{s(i),S(i)\}}} p(z',i,z)x(z',i) - w(z,u) + \sum_{v\in\Lambda} y(z,u,v) = 0 \quad (7)$$

$x, y, w \geq 0.$

The rate $x(z,i)$ or average number of transitions per unit time at which the process enters band i at the point z times $\mathbb{E}[T|(z,i)]$, the average time the process spends in the band when it enters at the point z, is simply the fraction of time the process spends in band i having entered it at point z. The *Scale Constraint* ensures that these fractions sum to one and thereby scales the formulation to one time unit. The *In-Conservation Constraint* for $(z,v) \in \mathcal{S}\times\Lambda$, ensures that the rate at which the process switches to the drift rate v at the point z is equal to the rate at which it enters bands with that drift rate at that point. Similarly, the *Out-Conservation Constraint* for $(z,u) \in \mathcal{S}\times\Lambda$ ensures that the rate at which the process switches from drift rate u at the point z is equal to the rate at which it reaches the boundary z of bands with that drift rate.

The variable $x(u)$ associated with the absorbing band $(u,-\infty,\infty)$ appears in the Scale Constraint with a coefficient of 1 because, if $x(u)$ is positive in a basic feasible solution, then $x(u)$ is the only positive variable in the solution and so we are only interested in the rate at which it accumulates cost. Thus, the linear programming formulation correctly models the fact that there are no transitions out of absorbing bands. These variables have no non-zero coefficients on the conservation constraints.

We define a family of functions called *relative value functions*, each of which satisfies an ordinary differential equation (ODE) known as the Poisson equation with certain boundary conditions.

For each basic feasible solution (x^*,y^*,w^*) to LP$(\Lambda,\mathcal{S})$, we define the corresponding *domain*

$$D(x^*,y^*,w^*) = \{(a,i) : i \text{ is an active band of } (x^*,y^*,w^*) \text{ and } a \in [s(i),S(i)]\}.$$

When considering a function f defined on a a domain like $D(x^*,y^*,w^*)$, we treat f as a family of functions $\{f(\cdot,i) : i \in \mathcal{I}\}$ each defined on the corresponding subset of *IR* and so, for example, use f' and f'' to represent derivatives with respect to the first argument.

A function $f : D(x^*, y^*, w^*) \mapsto IR$ is said to be a *relative value function* with respect to (x^*, y^*, w^*) if there is a scalar γ such that

$$\frac{\sigma^2}{2} f''(a,i) + u(i)f'(a,i) + c(u(i)) + ha = \gamma \tag{8}$$

for almost all $(a, i) \in D(x^*, y^*, w^*)$,

$$f(s(i), i) = f(s(i), \zeta(i)) + K(u(i), u(\zeta(i))) \tag{9}$$

for each active band i such that $s(i) \geq 0$,

$$f(S(i), i) = f(S(i), \tau(i)) + K(u(i), u(\tau(i))) \tag{10}$$

for each active band i such that $S(i) \leq \Theta$,

$$f'(0, i) = -U \quad \text{for each active band } i \text{ with } s(i) = -\infty \text{ and} \tag{11}$$

$$f'(\Theta, i) = M \quad \text{for each active band } i \text{ with } S(i) = \infty. \tag{12}$$

Corresponding to each basic feasible solution (x^*, y^*, w^*) to LP$(\Lambda, \mathcal{S})$ there is a relative value function that is unique up to an additive constant. In fact, we establish a correspondence between relative value functions and complementary dual solutions. Namely we show that

- for each complementary dual solution $(\gamma^*, \alpha^*, \beta^*)$ (where $\gamma^*, \alpha^*, \beta^*$ are the dual variables corresponding to the scale constraint, in-conservation and out-conservation constraints, respectively) and constant d the conditions

$$f(z, i) = \alpha^*(z, u(i)) + d \text{ for each } x^*(z, i) > 0$$
$$f(z, b(z, v)) = \beta^*(z, u) + d - K(u, v) \text{ for each } y^*(z, u, v) > 0,$$

  define a unique relative value function f with respect to (x^*, y^*, w^*),
- for each relative value function f with respect to (x^*, y^*, w^*) and constant d, any dual solution (γ, α, β) satisfying

$$\gamma = \text{AC}(x^*, y^*, w^*)$$
$$\alpha(z, u(i)) = f(z, i) + d \text{ for each } x^*(z, i) > 0$$
$$\beta(z, u) = f(z, b(z, v)) + K(u, v) + d \text{ for each } y^*(z, u, v) > 0,$$

  is a complementary dual solution.

Relying on the LP$(\Lambda, \mathcal{S})$ formulation, its dual and the relationship between the complementary dual solutions and the relative value functions we obtain several results. Theorem 1 ensures that there is an optimal solution to LP$(\Lambda, \mathcal{S})$ that is a deterministic non-overlapping control band policy.

Theorem 1. *For each finite set $\mathcal{S}$, LP$(\Lambda, \mathcal{S})$ admits an optimal policy that is a deterministic non-overlapping control band policy.*

In particular, we show that if we choose (x^*, y^*, w^*) to be a basic optimal solution to $\mathrm{LP}(\Lambda, \mathcal{S})$ with a smallest number of active bands, then (x^*, y^*, w^*) corresponds to a non-overlapping control band policy.

We observe that every basic feasible solution to $\mathrm{LP}(\Lambda, \mathcal{S})$ corresponds to a policy $\Phi \in \mathcal{P}(\mathcal{S})$, however not every policy in $\mathcal{P}(\mathcal{S})$ can be captured by a feasible solution of the linear program $\mathrm{LP}(\Lambda, \mathcal{S})$. Thus, a lower bound to the average cost of policies in $\mathcal{P}(\mathcal{S})$ provides a lower bound for the linear program $\mathrm{LP}(\Lambda, \mathcal{S})$ just as every feasible dual solution to the dual problem does. On the other hand a feasible dual solution does not necessarily provide a lower bound for the average cost of a policy in $\mathcal{P}(\mathcal{S})$. We construct a lower bound on the average cost of any policy in $\mathcal{P}(\mathcal{S})$ from an optimal solution to the dual of our linear program and so prove the stronger result, Theorem 2.

Theorem 2. *For each finite set Λ of drift rates and finite set $\mathcal{S}$ of candidate transition points, an optimal solution to $\mathrm{LP}(\Lambda, \mathcal{S})$ is an optimal policy for the $\mathcal{S}$-restricted Brownian control problem and so, the $\mathcal{S}$-restricted Brownian control problem admits an optimal policy that is a deterministic non-overlapping control band policy.*

Further, under mild assumptions on the cost coefficients, we can obtain deterministic non-overlapping control band policies that are asymptotically optimal for the unrestricted Brownian control problem by choosing consecutive points of $\mathcal{S}$ sufficiently close together. In particular, we show how to compute a lower bound on the average cost of any non-anticipating policy and, given $\epsilon > 0$, we show how to choose the points of $\mathcal{S}$ so that $\gamma(1-\epsilon)$ is such a lower bound, where γ is the average cost of an optimal solution to $\mathrm{LP}(\Lambda, \mathcal{S})$. These results are summarized in Proposition 1 and Theorem 3.

Note that when considering a function f defined on a subset of $I\!R \times \Lambda$ (or $I\!R \times \mathcal{I}$ for some finite index set $\mathcal{I}$), we treat f as a family of functions $\{f(\cdot, u) : u \in \Lambda\}$ ($\{f(\cdot, i) : i \in \mathcal{I}\}$) each defined on the corresponding subset of $I\!R$ and so, for example, use f' and f'' to represent derivatives with respect to the first argument.

Proposition 1. *Suppose that for each $u \in \Lambda$, $f(\cdot, u) : [0, \Theta] \to \mathbb{R}$ is a continuous function that can be written as the difference of two convex functions and*

i. is differentiable,
ii. has bounded first and second derivatives and
iii. has continuous second derivative

at all but a finite set of points. Further, suppose that for each $u \in \Lambda$ the function $f(\cdot, u)$ and the scalar γ satisfy:

$$\frac{\sigma^2}{2} f''(a, u) + u f'(a, u) + c(u) + ha \geq \gamma \text{ for a.a. } a \in [0, \Theta], \tag{13}$$

$$f(a, v) - f(a, u) \geq -K(u, v) \text{ for all } a \in [0, \Theta] \text{ and } v \in \Lambda, \tag{14}$$

$$f'(0, u) \geq -U, \tag{15}$$

$$f'(\Theta, u) \leq M, \tag{16}$$

$$\rho(a, u) \equiv f'(a+, u) - f'(a-, u) \geq 0 \text{ for all } a \in (0, \Theta). \tag{17}$$

Then $\gamma \leq \mathrm{AC}(\Phi)$ *for each policy* $\Phi \in \mathcal{P}$ *and each initial state* $(a, u) \in \mathcal{D}$*.*

Theorem 3. *Suppose* $h \geq 0$ *and* $U \geq \frac{c(u)}{u} \geq -M$ *for each* $u \in \Lambda$*, then for each* $\epsilon > 0$ *there is* $\delta > 0$ *such that if we choose the points of* $\mathcal{S}$ *in* $[0, \Theta]$ *so that consecutive points in* $\mathcal{S} \cup \{0, \Theta\}$ *are within* δ *of each other, then*

$$\gamma(1-\epsilon) \leq \inf_{\Phi \in \mathcal{P}} \mathrm{AC}(\Phi) \leq \min_{\Phi \in \mathcal{P}(\mathcal{S})} \mathrm{AC}(\Phi) = \gamma. \quad (18)$$

The proof of Theorems 1, 2 and 3 rely heavily on properties of complementary dual solutions and on the apparently new relationships we establish between complementary dual solutions and relative value functions.

4 Conclusion

In this paper we address a capacity management problem in a build-to-order setting and model it as a diffusion problem. An advantage of the diffusion problem is that it requires limited information about the process, namely, the mean and variance, and does not require any additional information like the distribution of demand which is usually hard to correctly identify and model. In solving the diffusion problem we define an $\mathcal{S}$-restricted Brownian control problem where the controller may change the drift rate only when the available to build level is equal to certain values. We initially restrict our attention to a class of policies called the control band policies that allow us to build an LP formulation. We show that the optimal policy to this LP is a deterministic non-overlapping control band policy and that this LP indeed solves the S-restricted problem by showing that the solution to the LP achieves a lower bound for the $\mathcal{S}$-restricted problem. Then we achieve asymptotic optimality by showing that under mild assumptions on cost parameters the average cost of an optimal solution to the S-restricted problem can be made arbitrarily close to this lower bound. Thus, this approach solves the unrestricted problem through different restricted versions of the problem and comes up with a simple policy. The simplicity of this policy greatly facilitates its application in industrial settings. The approach can easily be extended to settings that include the build-to-stock option. As the approach also isolates each cost component (and provides a simple expression to calculate) it helps the controller judge the impact of policy changes in terms of each cost component and its operational implications. One hopes to extend these results so that LP formulations of individual diffusion control problems can be knitted together and integrated into larger planning and control problems using traditional modeling techniques.

References

1. Ata, B., Harrison, J.M., Shepp, L.A.: Drift rate control of a Brownian processing system. The Annals of Applied Probability **15** (2005) 1145–1160

2. Avram, F., Karaesmen, F.: A method for computing double band policies for switching between two diffusions. Probability in Engineering Sciences **10** (1996) 569–590
3. Chernoff, H., Petkau, A.J.: Optimal control of a Brownian motion. SIAM Journal on Applied Mathematics **34** (1978) 717–731
4. Ghosh, A.P., Weerasinghe, A.P.: Optimal buffer size for a stochastic processing network in heavy traffic. Queueing Systems **55** (2007) 147–159
5. Liao, Y.C.: Switching and impulsive control of reflected diffusion. Applied Mathematical Optimization **11** (1984) 153–159
6. Perry, D., Bar-Lev, S.K.: A control of a Brownian storage system with two switchover drifts. Stochastic Analysis and Applications **7** (1989) 103–115
7. Rath, J.H.: The optimal policy for a controlled Brownian motion process. SIAM Journal on Applied Mathematics **32** (1977) 115–125
8. Manne, A.S.: Linear programming and sequential decisions. Management Science **6** (1960) 259–267
9. Feldman, R.M., Valdez-Flores, C.: Applied Probability and Stochastic processes. Pws Pub Co, Boston (1996)
10. Kushner, H.J., Dupuis, P.: Numerical methods for stochastic control problems in continuous time. Second edn. Volume 24 of Applications of Mathematics (New York). Springer-Verlag, New York (2001) Stochastic Modelling and Applied Probability.
11. Kumar, S., Muthuraman, K.: A numerical method for solving singular stochastic control problems. Operations Research **52** (2004) 563–582
12. Helmes, K., Stockbridge, R.: Numerical comparison of controls and verification of optimality for stochastic control problems. Journal of Optimization Theory and Applications **106** (2000) 107–127
13. Helmes, K., Stockbridge, R.: Determining the optimal control of singular stochastic processes using linear programming. In Ethier, S., Feng, J., Stockbridge, R., eds.: Markov Processes and Related Fields: A Festschrift in Honor of Thomas G. Kurtz. IMS lecture notes. (to appear)
14. Harrison, J.M.: Brownian motion and stochastic flow systems. Wiley, New York (1985)
15. Harrison, J.M., Williams, R.J.: Brownian models of open queueing networks with homogeneous customer populations. Stochastics **22** (1987) 77–115
16. Harrison, J.M., Williams, R.J.: Multidimensional reflected Brownian motions having exponential stationary distributions. Annals of Probability **15** (1987) 115–137