

Singular Control of (Reflected) Brownian Motion: A Computational Method Suitable for Queueing Applications

Baris Ata^{*1}, J. Michael Harrison^{†2}, and Nian Si^{‡1}

¹Booth School of Business, University of Chicago

²Graduate School of Business, Stanford University

Abstract

Motivated by applications in queueing theory, we consider a class of singular stochastic control problems whose state space is the d -dimensional positive orthant. The original problem is approximated by a drift control problem, to which we apply a recently developed computational method that is feasible for dimensions up to $d = 30$ or more. To show that nearly optimal solutions are obtainable using this method, we present computational results for a variety of examples, including queueing network examples that have appeared previously in the literature.

1 Introduction

In a recent paper [Ata et al., 2023] we described a new computational method for drift control of reflected Brownian motion, illustrating its use through applications to various test problems. Here we consider a class of singular stochastic control problems (see below for a discussion of this term) that are motivated by applications in queueing theory. Our approach is to approximate a problem in that class by a drift control problem of the type considered earlier, and we conjecture that, as a certain upper bound parameter becomes large, the solution obtained is nearly optimal for the original singular control problem. Numerical results will be presented here that give credence to that conjecture, but we do not attempt a rigorous proof.

Singular Stochastic Control. The modern theory of singular stochastic control was initiated by Beneš et al. [1980], whose work inspired subsequent research by Karatzas [1983], Harrison and Taksar [1983] and others in the 1980s, all of which focused on specially structured problem classes. In fact, no truly general definition of singular control has been advanced to this day, but the term is used in reference to problems having the three features specified in the following paragraph, where we denote by $W(t)$ the state of the system at time t .

^{*}baris.ata@chicagobooth.edu

[†]mike.harrison@stanford.edu

[‡]niansi@chicagobooth.edu

(a) In the absence of control, $W = \{W(t), t \geq 0\}$ evolves as a diffusion process whose state space S is a subset of d -dimensional Euclidean space. In all the work of which we are aware, the drift vector and covariance matrix of W are actually assumed to be constant (that is, to not depend on the current state), so the associated theory can be accurately described as singular control of Brownian motion, or of geometric Brownian motion in some financial applications, or of reflected Brownian motion in some queueing applications. (b) The system manager can enforce instantaneous displacements in the state process W , possibly including enforcement of jumps, in any of several fixed directions of control, subject to the requirement that W remain within the state space S . (c) The controller continuously incurs running costs, usually called holding costs in queueing theory applications, that are specified by a given function $h(W(t))$. He or she also incurs costs of control that are proportional to the sizes of the displacements enforced, with a different constant of proportionality for each direction of control.

A control policy takes the form of a p -dimensional cumulative control process $U = \{U(t), t \geq 0\}$ whose components are non-decreasing. Because the cumulative cost of control in direction i is proportional to $U_i(\cdot)$ by assumption, an optimal policy typically creates one or more endogenous reflecting barriers. For example, Harrison and Taksar [1983] considered a one-dimensional control problem (that is, $d = 1$) where W behaves as a (μ, σ) Brownian motion $X = \{X(t), t \geq 0\}$ in the absence of control, the holding cost function $h(\cdot)$ is convex, and the controller can displace $W(\cdot)$ either to the right at a cost of $r > 0$ per unit of displacement, or else to the left at a cost of $l > 0$ per unit of displacement.

Assuming the objective is to minimize expected discounted cost over an infinite planning horizon, an optimal policy enforces a lower reflecting barrier at some level a , and an upper reflecting barrier at some level $b > a$. That is, the two cumulative control processes $U_1(\cdot)$ and $U_2(\cdot)$, which correspond to rightward displacement and leftward displacement, respectively, increase in the minimal amounts necessary to keep $W(\cdot) = X(\cdot) + U_1(\cdot) - U_2(\cdot)$ confined to the interval $[a, b]$. Both $U_1(\cdot)$ and $U_2(\cdot)$ are almost surely continuous, but because the Brownian motion X has unbounded variation in every interval of positive length, each of them increases only on a set of time points having Lebesgue measure zero. In this sense, the optimal controls are almost surely singular with respect to Lebesgue measure; it is that feature of optimal controls, in a different but similar problem context, that led Beneš et al. [1980] to coin the term “singular control.”

To the best of our knowledge, the only previous works providing computational methods for singular control problems are those by Kushner and Martins [1991] and by Kumar and Muthuraman [2004]. In Section 7 of this paper we consider a two-dimensional problem and compare our results with those obtained using the Kushner-Martins method, finding that the two solutions are virtually indistinguishable. However, the methods propounded by both Kushner-Martins and Kumar-Muthuraman are limited to low-dimensional applications, because they rely on grid-based computations. To be specific, Kushner and Martins [1991] employ a Markov chain approximation, whereas Kumar and Muthuraman [2004] rely on the finite-element method for analysis of partial differential equations. In contrast, our simulation-based approximation scheme is suitable for high-dimensional problems,

two of which will be analyzed in Sections 4 and 8.

The rest of this paper. Section 2 describes in general terms the two closely related classes of problems to be treated here, and Section 3 explains our approach to such problems. Readers will see that our method can easily be extended to address other similar problems, some of which will be discussed later in Section 9. Still, the problem types described in Section 2 includes many important applications, and they are general enough to prove the viability of our approach via drift control approximations. In Section 4 our method is applied to a specially structured family of singular control problems that can be solved analytically. The numerical results presented there indicate that a high degree of accuracy is achievable with the method, and that it is computationally feasible for problems of dimension 30 or more.

Sections 5 through 8 are devoted to stochastic control problems associated with queueing network models, also called stochastic processing networks, cf. Dai and Harrison [2020]. More specifically, the problems considered in those sections involve dynamic control of networks in the “heavy traffic” parameter regime, for which singular control approximations have been derived in previous research. Analysis of such examples involves two steps: first, numerical solution of the approximating singular control problem; and second, translation of the singular control solution into an implementable policy for the original discrete-flow model. The second of those steps will be left as an open problem for one of our queueing network examples, because it involves what is essentially a separate research domain. However, we intend to continue analysis of that example in future work.

Section 9 contains concluding remarks on various topics, including the potential for generalization beyond the problems described in Section 2. Finally, there are several short appendices that provide added technical detail on subjects treated or referred to in the body of the paper.

The code used for the examples in this paper is available in <https://github.com/nian-si/SingularControl>.

2 Two classes of singular control problems to be treated

We consider a stochastic system with a d -dimensional *state process* $W = \{W(t), t \geq 0\}$ and a p -dimensional *control* $U = \{U(t), t \geq 0\}$. Two distinct problem formulations are to be treated, each of them generated from the following primitive stochastic element:

$$\begin{aligned} X = \{X(t), t \geq 0\} \text{ is a } d\text{-dimensional Brownian motion with drift vector } \xi, \\ \text{non-singular covariance matrix } A, \text{ and } X(0) = 0. \end{aligned} \tag{1}$$

Formulation with exogenous reflection at the boundary. In this problem W , U and a d -dimensional *boundary pushing process* Y jointly satisfy the following relationships and restrictions:

$$W(t) = w + X(t) + GU(t) + RY(t) \text{ for all } t \geq 0, \tag{2}$$

$$U(\cdot) \text{ is adapted to } X, \text{ right-continuous and non-decreasing with } U(0) \geq 0, \tag{3}$$

$$W(t) \geq 0 \text{ for all } t \geq 0, \quad (4)$$

$$Y = \{Y(t), t \geq 0\} \text{ is adapted to } X, \text{ continuous and non-decreasing with } Y(0) = 0, \text{ and} \quad (5)$$

$$Y_i(\cdot) \text{ only increases at times } t \geq 0 \text{ when } W_i(t) = 0 \text{ (} i = 1, \dots, d \text{)}. \quad (6)$$

Here $w \geq 0$ is a given initial state, G is a $d \times p$ matrix whose columns are interpreted as feasible *directions of control*, and R is a $d \times d$ *reflection matrix* that has the form

$$R = I - Q \text{ where } Q \text{ is non-negative and has spectral radius } \rho(Q) < 1. \quad (7)$$

Conditions (2) through (7) together define Y in terms of X and U , specifying that W is instantaneously reflected at the boundary of the orthant, and more specifically, that the direction of reflection from the boundary surface $W_i = 0$ is the i^{th} column of R . See below for further discussion.

Finally, we take as given a *holding cost function* $h : \mathbb{R}_+^d \rightarrow \mathbb{R}$, a vector $c \in \mathbb{R}_+^p$ of *control costs*, a vector $\pi \in \mathbb{R}_+^d$ of *penalty rates* associated with pushing at the boundary, and an interest rate $\gamma > 0$ for discounting. The system manager's objective is then to

$$\text{minimize } \mathbb{E} \left\{ \int_0^\infty e^{-\gamma t} [h(W(t))dt + c \cdot dU(t) + \pi \cdot dY(t)] \right\}. \quad (8)$$

The stochastic control problem (2)-(8) is identical to the one studied in our earlier paper [Ata et al., 2023] except that there we imposed two further restrictions. First, the vector of penalty rates was taken to be $\pi = 0$, but the computational method developed in [Ata et al., 2023] can easily be modified to cover the current more general case. Second, each component of the monotone control U was required to be absolutely continuous with a density bounded above by a given constant $b > 0$. It is the absence of this latter requirement (uniformly bounded rates of control), together with the linear cost of control assumed in (8), that distinguishes the current problem as one of singular control.

For each $j = 1, \dots, p$ and $t \geq 0$, one interprets $U_j(t)$ as the cumulative amount of displacement effected over the time interval $[0, t]$ in the j^{th} available direction of control (that is, in the direction specified by the j^{th} column of G), and one interprets $c_j U_j(t)$ as the associated cumulative cost. Condition (3) allows instantaneous displacements in any of the feasible directions of control, represented by jumps in the corresponding components of U , including jumps in one or more of those directions at $t = 0$, represented by positive values for different components of $U(0)$.

Our use of the letter W to denote system state, as opposed to the letter Z that was used for that purpose in Ata et al. [2023], is motivated by applications in queueing theory (see Sections 5 through 8), where it is mnemonic for *workload*. In general, those queueing theoretic applications give rise to workload processes whose state space is a d -dimensional polyhedral cone, as opposed to the orthant that is specified by (4) as the state space for our control problem here. However, a polyhedral cone can always be transformed to an orthant by a simple change of variables, so one may say that here we present our control problem in a convenient *canonical form*, with the goal of simplifying the exposition.

A square matrix R of the form (7) is called a *Minkowski matrix* in linear algebra (or just *M-matrix* for brevity). It is non-singular, and its inverse is given by the Neumann expansion $R^{-1} = I + Q + Q^2 + \dots$. A process U that satisfies conditions (2) through (6) will be called a *feasible control*, and by assuming that R is an M-matrix we ensure the existence of at least one feasible control, as follows. If one simply sets $U(\cdot) = 0$, then a result by Harrison and Reiman [1981] shows that there exist d -dimensional processes W and Y that jointly satisfy (2) through (6). Moreover, W and Y are unique in the pathwise sense, and Y is adapted to X . Thus the control $U(\cdot) = 0$ is feasible. The resulting state process W is called a *reflected Brownian motion* (RBM) with reflection matrix R .

Formulation without exogenous reflection at the boundary. In our second problem formulation, the assumption of exogenous reflection at the boundary is removed, so system dynamics are specified as follows:

$$W(t) = w + X(t) + GU(t) \text{ for all } t \geq 0, \quad (9)$$

$$U(\cdot) \text{ is adapted to } X, \text{ right-continuous, and non-decreasing with } U(0) \geq 0, \text{ and} \quad (10)$$

$$W(t) \geq 0 \text{ for all } t \geq 0. \quad (11)$$

As before, G is a $d \times p$ matrix whose columns are interpreted as feasible directions of control, but now we assume that $p \geq d$ and

$$\text{the first } d \text{ columns of } G \text{ form an M-matrix that we denote hereafter as } R. \quad (12)$$

Finally, given a holding cost function $h(\cdot)$ and control cost vector c as before, the system manager's objective is to

$$\text{minimize } \mathbb{E} \left\{ \int_0^\infty e^{-\gamma t} [h(W(t))dt + c \cdot dU(t)] \right\}. \quad (13)$$

Assumption (12) is satisfied in most queueing applications, as we shall see in Sections 5 through 8, and it guarantees the existence of a feasible control in the following obvious manner. First, given the underlying Brownian motion X , let W be the corresponding RBM and Y its associated boundary pushing process, both constructed from X as in Harrison and Reiman [1981] using R as the reflection matrix. Then define a p -dimensional control U by taking $U_i = Y_i$ for $i = 1, \dots, d$ and $U_i = 0$ for $i = d + 1, \dots, p$. This specific control U is continuous but in general not absolutely continuous. Of course, a feasible control that minimizes the objective in (13) may be very different in character.

3 Approximation by a drift control problem

This section explains our computational approach to each of the problem classes introduced in Section 2, but taking them in reverse order. Readers will see that all of our examples can be viewed as applications of the formulation *without* exogenous reflection at the boundary, which is one reason for treating it first.

Formulation without exogenous reflection at the boundary. Let us first consider the sin-

gular control problem *without* exogenous reflection at the boundary, defined by system relationships (9)-(11), our crucial assumption (12) on the control matrix G , and the objective (13). To approximate this problem, we restrict attention to controls U that satisfy conditions (14)-(18) below, where W is again defined in terms of X and U via (9), W must again be confined to the non-negative orthant as in (11), and the objective (13) remains unchanged. The first restriction on controls U is that they have the following form: for each $t \geq 0$,

$$U_i(t) = \int_0^t \theta_i(s)ds + Y_i(t) \quad \text{for } i = 1, \dots, d \text{ and} \quad (14)$$

$$U_j(t) = \int_0^t \theta_j(s)ds \quad \text{for } j = d+1, \dots, p, \quad (15)$$

where $\theta(\cdot)$ is a p -dimensional *drift control*, and $Y(\cdot)$ is a corresponding d -dimensional *boundary control* to be specified shortly. Second, it is required that $\theta(\cdot)$ be adapted to X and satisfy

$$0 \leq \theta_i(\cdot) \leq b \text{ for each } i = 1, \dots, p, \text{ where } b > 0 \text{ is a given constant.} \quad (16)$$

Third, $Y(\cdot)$ is defined in terms of X and θ by the following relationships, which are identical to (5) and (6):

$$Y(\cdot) \text{ is adapted to } X, \text{ continuous and non-decreasing with } Y(0) = 0, \text{ and} \quad (17)$$

$$Y_i(\cdot) \text{ only increases at times } t \geq 0 \text{ when } W_i(t) = 0 \text{ (} i = 1, \dots, d \text{)}. \quad (18)$$

By combining (14) and (15) with our assumption (12) about the form of the control matrix G , one sees that the main system equation (9) can be rewritten as follows when controls are of the restrictive form considered here:

$$W(t) = w + X(t) + G \int_0^t \theta(s)ds + RY(t) \text{ for all } t \geq 0. \quad (19)$$

As noted earlier, conditions (17) and (18) are the standard means of specifying instantaneous reflection at the boundary of the orthant. More specifically, the direction of reflection from the boundary surface $W_i = 0$ is the i^{th} column of R , or equivalently, the i^{th} column of G ($i = 1, \dots, d$). One interprets $Y_i(t)$ as the cumulative amount of displacement effected in that direction over the time interval $[0, t]$ in order to prevent W from exiting the positive orthant. Note that the objective (13) associates a cost of c_i with each unit of increase in $Y_i(\cdot)$, just as it continuously accumulates drift-related costs at rate $c_i \theta_i$ per time unit ($i = 1, \dots, d$).

In our current setting, the boundary controls are redundant in an obvious sense: the i^{th} column of R is identical to the i^{th} column of G , so “reflection” at the boundary surface $W_i = 0$ displaces the system state W in the same direction as does the i^{th} component of the system manager’s drift control $\theta(\cdot)$, and we associate the same linear cost rate with increases in $Y_i(\cdot)$ as we do with positive values of $\theta_i(\cdot)$. The only reason for including reflection at the boundary as a “backup capability”

in our formulation is to ensure that feasibility can be maintained despite the upper bound b on allowable drift rates. In the limit as $b \uparrow \infty$, one expects that backup capability to be increasingly irrelevant. For example, if displacement by means of $Y_i(\cdot)$ were made even slightly more expensive than displacement by means of $\theta_i(\cdot)$, one would expect increases in $Y_i(\cdot)$ under an optimal policy to be rare and small as b becomes large.

The approximating drift control problem specified immediately above is of the form considered in our earlier study [Ata et al., 2023], where a computational method for its solution was developed and illustrated. As in that study, attention will be restricted here to stationary Markov policies, by which we mean that

$$\theta(t) = u(W(t)), t \geq 0, \text{ for some measurable policy function } u : \mathbb{R}_+^d \rightarrow [0, b]^p. \quad (20)$$

Formulation with exogenous reflection at the boundary. Secondly, let us consider the problem *with* exogenous reflection at the boundary, defined earlier by the system relationships (2)-(6), assumption (7) on the form of the reflection matrix R , and the objective (8). In this case we restrict attention to controls U of the form

$$U(t) = \int_0^t \theta(s) ds, \quad (21)$$

where $\theta(\cdot)$ is a p -dimensional drift control adapted to X and satisfying

$$0 \leq \theta_i(\cdot) \leq b \text{ for each } i = 1, \dots, p, \text{ where } b > 0 \text{ is a given constant.} \quad (22)$$

Given a control U of this form, we define a corresponding boundary pushing process Y and state process W via (17)-(19) as before. Finally, the objective for our approximating drift control formulation is again given by (8), which involves a vector π of penalty rates associated with the various components of Y , separate from the cost rate vector c associated with the endogenous control U .

4 Comparison with known solutions

Let us consider a one-dimensional ($d = 1$) singular control problem with $p = 2$ directions of control (up and down, or right and left), initially framing the problem as one without exogenous reflection at the boundary. To be specific, let us assume the following problem data:

$$\xi = 0, A = [1], G = [1, -1], c = (0, 1), h(w) = 2w \text{ and } \gamma = 0.1.$$

That is, we solve the following singular control problem:

$$\begin{aligned} & \min_{U(\cdot)} \mathbb{E} \left\{ \int_0^\infty e^{-0.1t} [2W(t) + dU_2(t)] \right\} \\ & \text{subject to} \end{aligned}$$

$W(t) = w + X(t) + U_1(t) - U_2(t)$, $t \geq 0$,
 $X = \{X(t), t \geq 0\}$ is a one-dimensional standard Brownian motion,
 $U(\cdot)$ is adapted to X , right-continuous and non-decreasing with $U(0) \geq 0$, and
 $W(t) \geq 0$ for all $t \geq 0$.

Because the holding cost function $h(\cdot)$ is positive and increasing, it is obvious that control component U_1 , which pushes the state process W upward or rightward, should be minimal. Equivalently stated, U_1 should be structured to enforce a lower reflecting barrier at zero. Thus the problem under discussion can be reframed as one of singular control *with* exogenous reflection at the boundary, and it is analyzed (for general problem data) in those terms in Appendix D. There we show that the optimal choice for U_2 enforces an upper reflecting barrier at a threshold value $w^* > 0$, so the optimal policy confines W to a closed interval $[0, w^*]$, and we provide a formula for w^* .

Comparing solutions with different values for the upper bound b . Now consider the drift control approximation (with exogenous reflection at the boundary) discussed in Section 2, specialized to the one-dimensional problem described above. The analytic solution of that problem was derived in Appendix D.2 of our previous paper Ata et al. [2023]. The optimal policy is of a bang-bang type, featuring a specific threshold level: the optimal control employs a downward drift at the maximum allowable rate b when $W(\cdot)$ is above the threshold, and employs zero drift otherwise. Table 1 compares the optimal threshold values, and the objective values $V(0)$ achieved, for the drift control approximations with various values of b and for our "exact" singular control formulation. It is notable that even with the very tight bound $b = 5$, the approximation error expressed in terms of the system manager's objective is less than 1%.

b	2	5	10	20	100	Singular
Threshold	0.52	0.63	0.67	0.70	0.72	0.72
$V(0)$	14.73	14.09	14.00	13.97	13.96	13.96

Table 1: Analytic solutions with different upper bounds b of the one-dimensional problem.

A decomposable multi-dimensional singular control problem. Let us now consider a d -dimensional problem of the form specified in Section 2 (*without* exogenous reflection at the boundary) and the following data: control dimension $p = 2d$, drift vector $\xi = 0$, covariance matrix $A = I_{d \times d}$, control cost vector $c = (0, 0, \dots, 0, 1, 1, \dots, 1)$, holding cost function $h(w) = 2w_1 + \dots + 2w_d$, interest rate $\gamma = 0.1$, and control matrix

$$G = \begin{bmatrix} 1 & & & -1 & & \\ & 1 & & & -1 & \\ & & \ddots & & & \ddots \\ & & & 1 & & -1 \end{bmatrix}$$

This means that the d components of W are controlled independently, and the control problem for

each individual component is the one-dimensional singular control problem discussed immediately above. This can be interpreted as a heavy-traffic approximation for a queueing network model where d single-server queues operate independently and in parallel (see Figure 1); a system manager can shut off the input to any one of the queues at any time, incurring a cost of 1 per unit of time that input is denied, and continuously incurs a cost of 2 per job held in the queue or in service.

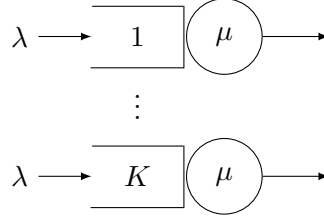


Figure 1: A decomposable parallel-server queueing network.

We solved this 30-dimensional problem using our computational method with an upper bound value $b = 10$, making no use of the problem's decomposability, and then we estimated the objective value $V(0)$ for our computed solution using Monte Carlo simulation. (That is, we repeatedly simulated the evolution of the Brownian system model with starting state $w = 0$ and operating under our computed optimal policy, recording the average objective value and an associated confidence interval.) The result is reported in Table 2, where we also record the simulated performance of the known analytic solution. Of course, the performance figures reported in Table 2 are subject to simulation and discretization errors. The run-time for our method is about twenty-four hours in this case using a 20-CPU core computer.

Analytic optimal solution	Our method
417.4 ± 0.2	416.8 ± 0.2

Table 2: Simulation performance for the 30-dimensional parallel-server network.

5 Tandem queues example

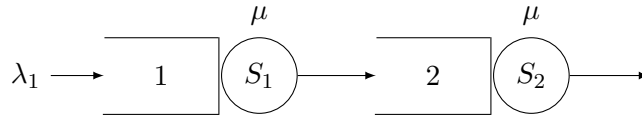


Figure 2: A network of tandem queues.

Figure 2 pictures a network model with two job classes and two single-server stations. Class 1 jobs arrive from outside the system according to a Poisson process with rate λ , and they are served at station 1 on a first-in-first-out (FIFO) basis. After completing service at station 1, they become class 2 jobs and proceed to station 2, where they are again served on a FIFO basis. Class 2 jobs

leave the system after completing service. For jobs of each class there is a similarly numbered buffer (represented by an open-ended rectangle in Figure 2) where jobs reside both before and during their service.

For simplicity, we assume that service times at each station are i.i.d. exponential random variables with mean $m < \lambda^{-1}$, and furthermore, any service can be interrupted at any time and resumed later without any efficiency loss. We denote by $Q_k(t)$ the number of class k jobs in the system at time t ($k = 1, 2$), calling this a "queue length process," and when we speak of the "system state" at time t , that means the pair $Q(t) = (Q_1(t), Q_2(t))$. There is a holding cost rate of h_k per class k job per unit of time, so the instantaneous holding cost rate at time t is $h_1 Q_1(t) + h_2 Q_2(t)$.

We shall assume the specific numerical values $\lambda = 0.95$ and $m = 1$, with holding costs $h_1 = 1$ and $h_2 = 2$, and with interest rate $r = 0.01$ for discounting. The system manager's control policy specifies whether the two servers are to be working or idle in each system state. If one is given holding costs such that $h_1 \geq h_2 > 0$, then an optimal policy is to have each server keep working so long as there are customers present in that server's buffer, but with the data assumed here, it is more expensive to hold class 2 than class 1, so there is a potential motivation to idle server 1 even when there are customers present in its buffer. The system manager seeks a control policy to minimize total expected discounted holding costs over an infinite planning horizon.

Following the development in Harrison [1988], we imagine this network model as part of a sequence indexed by $n = 1, 2, \dots$ that approaches a heavy traffic limit. In the current context, the assumption of "heavy traffic" can be made precise as follows: there exists a large integer n such that $\sqrt{n}(\lambda - m^{-1})$ is of moderate absolute value. With the system parameters specified above, a plausible value of the associated sequence index is $n = 400$, which gives $|\sqrt{n}(\lambda - m^{-1})| = 1$. Using that value of n , we define a scaled queue length process $Z(t) = (Z_1(t), Z_2(t))$ as follows:

$$Z(t) = n^{-1/2}Q(nt) \text{ for } t \geq 0. \quad (23)$$

In a stream of work including Harrison [1988] and Harrison and Van Mieghem [1997], the following two-stage scheme was developed for approximating a queueing control problem in the heavy traffic parameter regime. First, the original problem is approximated by a *Brownian control problem* (BCP) whose state descriptor $Z(t)$ corresponds to the scaled queue length process defined above. Second, the BCP is replaced by a *reduced Brownian control problem*, also called the *equivalent workload formulation* (EWF), whose state descriptor $W(t)$ corresponds to a certain linear transformation of Z and is called a *workload process*. For the simple example considered in this section, W is actually identical to Z , and there is no difference between the original BCP and its EWF. To be specific, the BCP is a singular control problem of the form specified in Section 2 (the formulation *without* exogenous reflection at the boundary), with drift vector

$$\xi = (\sqrt{n}(\lambda - m^{-1}), 0) = (-1, 0),$$

covariance matrix

$$A = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix},$$

control matrix

$$G = \begin{bmatrix} 1 & 0 \\ -1 & 1 \end{bmatrix},$$

holding cost function

$$h(w) = h_1 w_1 + h_2 w_2 = w_1 + 2w_2,$$

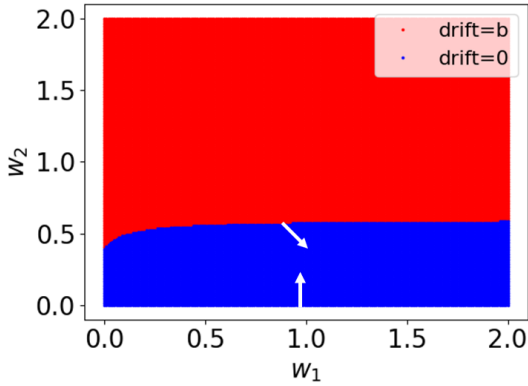
control cost vector $c = 0$, and interest rate for discounting

$$\gamma = nr = 4.$$

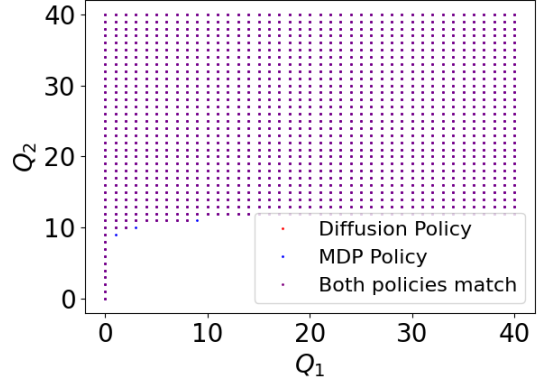
Applying our computational method to this singular control problem, we set the upper bound on the feasible drift rates at $b = 20$. The resulting optimal policy has the bang-bang form shown in the left-hand panel of Figure 3: the control component U_1 (corresponding to idleness of server 1) increases at the maximum allowable rate in the red region of the state space, and it increases not at all in the complementary blue region, except when the vertical axis is hit; on the other hand, U_2 increases only when W reaches the horizontal axis. In our singular control formulation, increases in U_1 produce an instantaneous displacement downward and to the right throughout the red region, whereas increases in U_2 produce instantaneous reflection in the vertical direction from the boundary $W_2 = 0$.

To interpret this solution in our original problem context, readers should recall that state (W_1, W_2) for the Brownian control problem corresponds to system state $(Q_1, Q_2) = (20W_1, 20W_2)$ in the original tandem queueing network. The obvious interpretation of our computed solution is the following: server 2 is never idled except when it has no work to do (that is, except when buffer 2 is empty), but server 1 is idled if either it has no work to do or W lies in the region colored red in Figure 3; the exact computational meaning of this description will be spelled out in Appendix C.

In addition to solving the singular control problem, we have also solved the Markov decision problem (MDP) associated with the original tandem queueing network. The right-hand panel of Figure 3 displays the optimal policies derived from both the exact MDP formulation and the approximating singular control formulation; the latter is referred to as the "diffusion policy" in the figure. For the MDP solution, blue dots signify states in which server 1 is idled, and for the diffusion policy, red dots signify such states. Figure 3 shows that the two policies are strikingly close. Also, Table 3 reports the simulated performance (that is, the average present value of all holding costs incurred) for the two policies applied in our original queueing network, assuming both buffers are initially empty, along with standard errors. For comparison, Table 3 also shows the performance of the policy that never idles either server except when its buffer is empty. In our simulations, one sees that the proposed policy derived from the singular control approximation actually performs as well as the optimal policy from the MDP formulation.



(a) Singular control policy



(b) Two policies for tandem queues

Figure 3: Optimal policies for the singular control and MDP formulations.

Never idle	MDP	Diffusion
1780 ± 1.0	1703 ± 0.9	1703 ± 0.9

Table 3: Simulated performance for the tandem queues example.

6 The criss-cross network with various cost structures

As shown in Figure 4, the criss-cross network consists of two single-server stations serving jobs of three different classes. Jobs arrive to classes 1 and 2 according to independent Poisson processes with rates λ_1 and λ_2 , respectively, and both of those classes are served at station 1. After completing service, class 1 jobs leave the system, whereas class 2 jobs make a transition to class 3 and are then served at station 2. As in Section 5, we suppose that jobs of each class reside in a similarly numbered buffer (represented by an open-ended rectangle in Figure 4) both before and during their service.

Class k jobs have exponentially distributed service times with mean m_k ($k = 1, 2, 3$), and the three service time sequences are mutually independent. Also, for maximum simplicity, we assume again that any service can be interrupted at any time and resumed later without any efficiency loss. A system manager decides, at each point in time, which job class to serve at each station, including the possibility of idleness for either server. Because we assume a positive holding cost rate for each class, and because only class 3 jobs are served at station 2, an optimal policy will obviously keep server 2 busy serving class 3 whenever buffer 3 is non-empty. Thus the problem boils down to deciding, at each point in time, whether to serve a class 1 job at station 1, serve a class 2 job there, or idle server 1.

The criss-cross network was introduced by Harrison and Wein [1989], who restricted attention to one particular set of cost parameters and focused on the heavy traffic regime. Martins et al. [1996] revisited the same model, also focused on the heavy traffic regime, but allowed more general cost parameters. The latter authors identified five mutually exclusive and collectively exhaustive cases

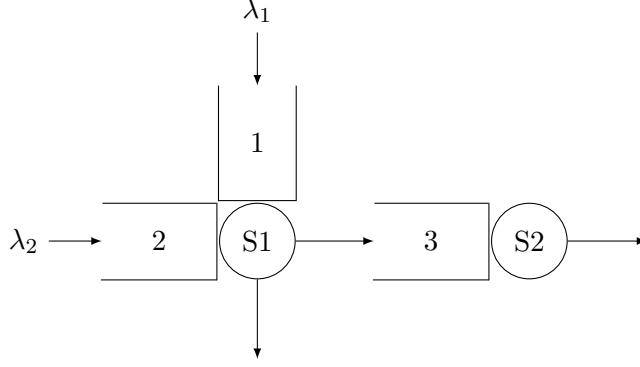


Figure 4: The criss-cross network.

with respect to cost parameters, and they offered conjectures about the form of an optimal policy in each case. More will be said about this early work below.

As in Section 5, we denote by $Q_k(t)$ the number of class k jobs in the system at time $t \geq 0$, or equivalently, the content of buffer k at time t ($k = 1, 2, 3$). Defining the system state $Q(t) = (Q_1(t), Q_2(t), Q_3(t))$, and denoting by $h = (h_1, h_2, h_3)$ the vector of positive holding cost parameters, the instantaneous cost rate incurred by the system manager at time t is given by the inner product $h \cdot Q(t)$. Harrison and Wein [1989] took the holding cost parameters to be $h_1 = h_2 = h_3 = 1$, which means that the instantaneous cost rate in every state equals the total number of jobs in the system. Martins et al. [1996] allowed any combination of positive holding cost parameters, and they identified the following cases:

$$\text{Case I: } (h_2 - h_3)/m_2 \geq h_1/m_1 \quad (24)$$

$$\text{Case II: } (h_2 - h_3)/m_2 < h_1/m_1 \quad (25)$$

$$\text{Case IIA: (25) plus } h_2 \geq h_3 \text{ and } h_2/m_2 \geq h_1/m_1 \quad (26)$$

$$\text{Case IIB: (25) plus } h_2 < h_3 \text{ and } h_2/m_2 \geq h_1/m_1 \quad (27)$$

$$\text{Case IIC: (25) plus } h_2 \geq h_3 \text{ and } h_2/m_2 < h_1/m_1 \quad (28)$$

$$\text{Case IID: (25) plus } h_2 < h_3 \text{ and } h_2/m_2 < h_1/m_1. \quad (29)$$

The simplest of these is Case I, which was studied by Chen et al. [1994] and will not be discussed further here. Rather, we focus attention on Case II and numerically solve examples covering its four sub-cases.

Prior to that analysis, a few more words about existing literature are appropriate. The specific cost structure studied by Harrison and Wein [1989] is an example of Case IIA. Those authors proposed an effective policy for that case, and Martins et al. [1996] proved the asymptotic optimality of that policy; also see Budhiraja and Ghosh [2005]. Martins et al. [1996] further made conjectures about the remaining sub-cases, two of which (Cases IIB and IIC) were rigorously justified in later papers by Budhiraja and Ross [2008] and Budhiraja et al. [2017].

Proceeding now to the analysis of numerical examples, we shall analyze the four holding cost combinations specified in Table 4. The stochastic parameters in all four numerical examples will be $\lambda_1 = \lambda_2 = 0.95$, $m_1 = m_2 = 0.5$ and $m_3 = 1$. The load factor for each server is then 0.95, meaning that 95 percent of each server's time is required to process all arrivals. With these data we are in the heavy traffic regime, and as in Section 5, a reasonable choice of the sequence index or scaling parameter is $n = 400$. Finally, the interest rate for discounting is again taken to be $r = 0.01$.

	h_1	h_2	h_3
IIA	1	1	1
IIB	1	1	1.5
IIC	1.5	1	1
IID	1.5	1	1.5

Table 4: Holding cost parameters for different cases.

Reference was made in Section 5 to a two-stage procedure for deriving a heavy traffic approximation to a queueing network control problem. For the criss-cross network, that analysis was undertaken by Harrison and Wein [1989], as follows. First, a Brownian control problem (BCP) is formulated whose state vector $Z(t)$ corresponds to the three-dimensional scaled queue length process defined via (23) as in Section 5. Then an equivalent workload formulation (EWF) is derived whose state vector $W(t)$ is defined for the criss-cross network as

$$W(t) = MZ(t) \text{ for } t \geq 0, \quad (30)$$

where

$$M = \begin{bmatrix} 0.5 & 0.5 & 0 \\ 0 & 1 & 1 \end{bmatrix}.$$

Obviously, the process W inherits from Z the scaling of time by a factor of n and the scaling of queue lengths by a factor of \sqrt{n} . We call M a *workload profile matrix*, interpreting M_{kj} as the expected amount of time that server k must spend to complete the processing of a job currently in class j . Thus, except for the scale factors referred to above, one interprets $W_k(t)$ as the expected amount of time required from server k to complete the processing of all jobs present anywhere in the system at time t ($k = 1, 2$). The control process $U(t)$ for the EWF is also two-dimensional, with $U_k(t)$ interpreted as a scaled version of the cumulative idleness experienced by server k up to time t ($k = 1, 2$).

The EWF derived by Harrison and Wein [1989] for the criss-cross network is a singular control problem of the form (9)-(13) specified in Section 2 (again we use the formulation *without* exogenous reflection at the boundary), with state dimension $d = 2$, control dimension $p = 2$, drift vector

$$\xi = \sqrt{n} (\lambda_1 m_1 + \lambda_2 m_2 - 1, \lambda_2 m_3 - 1) = (-1, -1),$$

covariance matrix

$$A = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 2 \end{bmatrix},$$

control matrix

$$G = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

control cost vector $c = 0$, and interest rate for discounting

$$\gamma = nr = 4.$$

Finally, the holding cost function $h(w)$ for the EWF is

$$h(w) = \min\{h \cdot z : Mz = w, z \geq 0\}, \quad (31)$$

which specializes as in Table 5 for the stochastic parameters assumed earlier and the holding cost parameters specified in Table 4.

	If $w_2 \geq 2w_1$	If $w_2 < 2w_1$
IIA	w_2	$2w_1$
IIB	$-w_1 + 1.5w_2$	$2w_1$
IIC	w_2	$3w_1 - 0.5w_2$
IID	$-w_1 + 1.5w_2$	$3w_1 - 0.5w_2$

Table 5: Cost function $h(w_1, w_2)$ for numerical values in Table 4.

An intuitive explanation of (31) is as follows. In the heavy traffic regime, a system manager can navigate quickly (that is, instantaneously in the heavy traffic limit) and costlessly between any two (scaled) queue length vectors z that yield the same (scaled) workload vector w . In our criss-cross model, such navigation is accomplished by adjusting service priorities at station 1: giving priority there to class 1 drives z_1 quickly down and z_2 quickly up in equal amounts, and vice-versa, without affecting w . Thus, the system manager can dynamically adjust priorities at station 1 so as to hold, at each point in time t , the least costly queue length vector $Z(t)$ that is consistent with the workload vector $W(t)$ that then prevails.

To review, heavy traffic analysis of the criss-cross network begins by formulating a three dimensional Brownian control problem (BCP) whose state vector $Z(t)$ is a scaled version of the model's three-dimensional queue length process $Q(t)$. Then the BCP is reduced to a two-dimensional equivalent workload formulation (EWF) whose state vector $W(t)$ is defined in terms of $Z(t)$ via (30), and the non-linear cost function for the EWF is given by (31). Actually, there are four different versions of the criss-cross EWF, using the four different cost structures specified in Table 5. We have applied our computational method to all four of those singular control problems, setting the algorithm's upper bound on drift rates at $b = 20$ as in our previous analysis of tandem queues in Section 5.

The resulting solutions are displayed graphically in Figure 5, where displacement to the right corresponds to idleness of server 1, and upward displacement corresponds to idleness of server 2. In Case IIA, each server is idled only when there is no work for it to do anywhere in the system. In the other three cases, our solution applies maximal drift to the right, interpreted as idleness of server 1, in the upper red region if there is one, and applies maximal upward drift, interpreted as idleness of server 2, in the lower red region if there is one. On the other hand, the computed solution applies no drift in the blue region except at the axes, which is interpreted to mean that each server continues to work at full capacity in those regions except when there is no work for that server anywhere in the system.

To understand why and how these idleness prescriptions are to be accomplished, one must consider the rest of the computed solution, that is, the scaled queue length vector z in which the scaled workload w is to be held. Equations (8.5a) and (8.5b) of Martins et al. [1996] show that the minimum in (31) is achieved in all four of our cases by

$$z^*(w) = ((2w_1 - w_2)^+, 2w_1 \wedge w_2, (w_2 - 2w_1)^+). \quad (32)$$

Superficially, this appears to say, in Case IIA for example, that buffer 3 should remain empty whenever W falls below the dotted white line in Figure 5, even though server 2 is to be idled only when W reaches the horizontal axis. To resolve this seeming contradiction, Harrison and Wein [1989] interpreted (32) to mean the following: class 1 (the class that exits the system when its one service is completed) should be given priority at station 1 whenever $Q_1(t) > 0$ and $Q_3(t) > s$ (mnemonic for *safety stock*), where $s > 0$ is a tuning parameter to be optimized via simulation; otherwise Class 2 should be given priority. The intuitive argument supporting this policy is that granting default priority to class 1 greedily reduces the instantaneous cost rate incurred by the system manager, but by taking s large enough one can still protect server 2 from starvation.

Martins et al. [1996] generalized the policy proposed by Harrison and Wein [1989] to cover all four cases as follows. First, idle server 1 whenever W reaches the vertical axis or falls in the upper red region in Figure 5. Second, idle server 2 whenever W reaches the horizontal axis or falls in the lower red region in Figure 5; this happens through prioritizing class 1 at server 1, which starves server 2. Third, give class 1 priority at station 1 if W falls in the blue region and $Q_3 > s$; otherwise, give priority to class 2. Here again, $s > 0$ is a tuning parameter to be optimized via simulation, and readers may consult Appendix C for the exact computational meaning of the verbal policy description given here. Hereafter, our policy derived from the singular control approximation will be referred to as the *diffusion policy* for the criss-cross network.

For comparison, we solved the associated MDP numerically to obtain an exact optimal policy. Figure 6 displays the optimal MDP policy. Motivated by (32), one expects $Q_1(t) \approx 0$ whenever the workload vector is above the dotted white line in Figure 5. Thus, we restrict attention to the case $Q_1 = 0$ in panels (a) and (c) of Figure 6 and display the optimal action as a function of Q_2 and Q_3 . For the MDP policy, blue dots signify states in which server 1 is idled, and for our diffusion policy derived from the approximating singular control problem (see Algorithm 1 in Appendix C), red dots

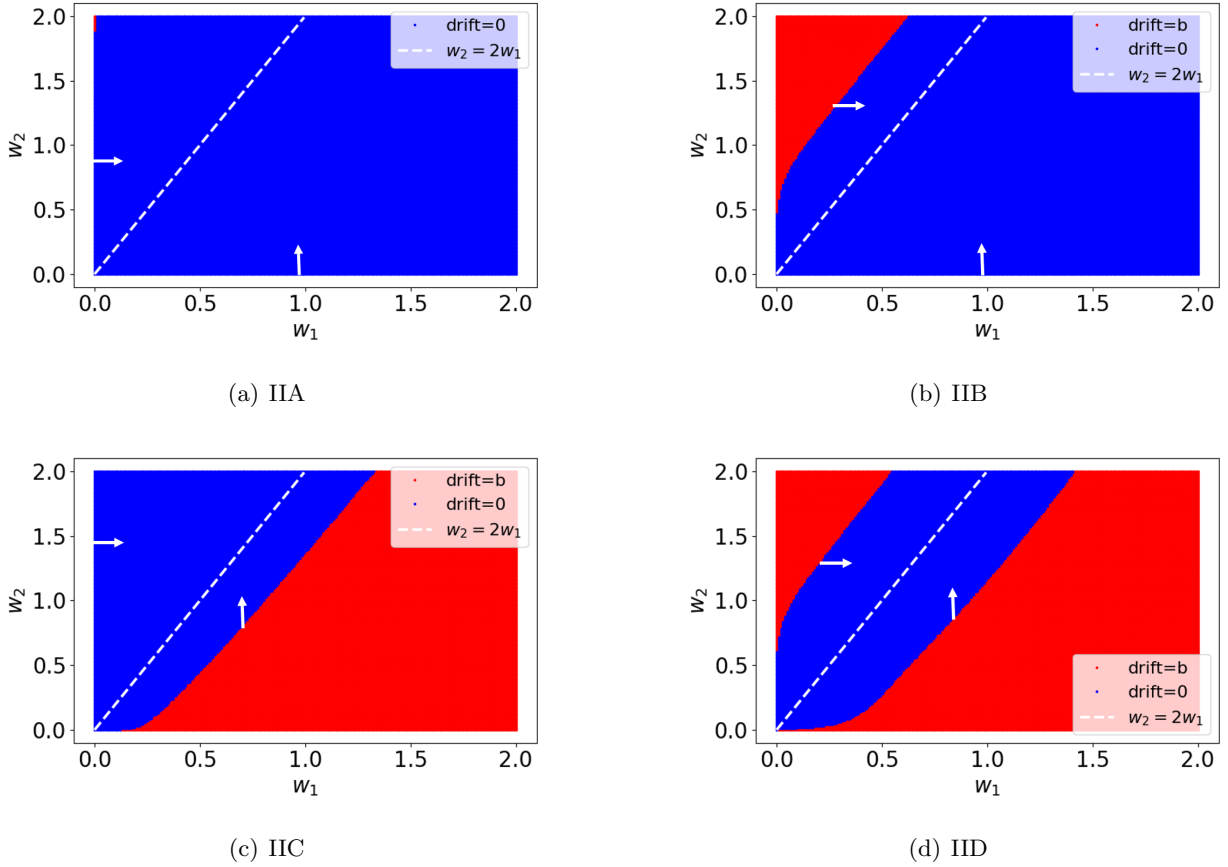


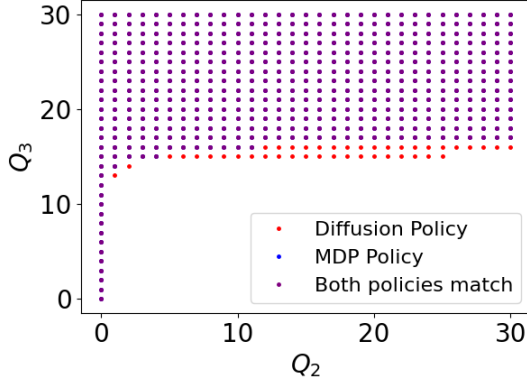
Figure 5: Graphical representation of the diffusion policies derived via solution of the approximating singular control problem.

signify such states. Similarly, one expects $Q_3(t) \approx 0$ whenever the workload vector falls below the dotted white line. Therefore, we restrict attention to the case $Q_3 = 0$ in panels (b) and (d) of Figure 6 and display the optimal action as a function of Q_1 and Q_2 . For the MDP policy, blue dots signify states in which server 1 prioritizes class 1, which starves server 2 and causes it to idle, and for the diffusion policy, red dots signify such states. Figure 6 shows that the two policies are very similar.

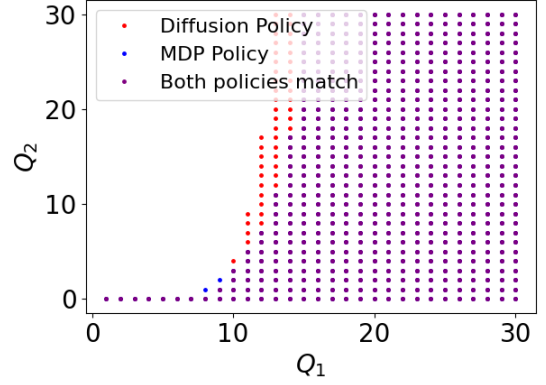
Cases	Static priority	MDP	Our proposed policy
IIA	1765 ± 1.1	1488 ± 0.9	1491 ± 0.9 ($s^* = 2$)
IIB	2084 ± 1.1	1659 ± 1.0	1665 ± 1.0 ($s^* = 2$)
IIC	1812 ± 1.0	1684 ± 1.1	1690 ± 1.0 ($s^* = 1$)
IID	2134 ± 1.1	1872 ± 1.0	1880 ± 1.1 ($s^* = 1$)

Table 6: Simulated performance of three policies for cases IIA - IID.

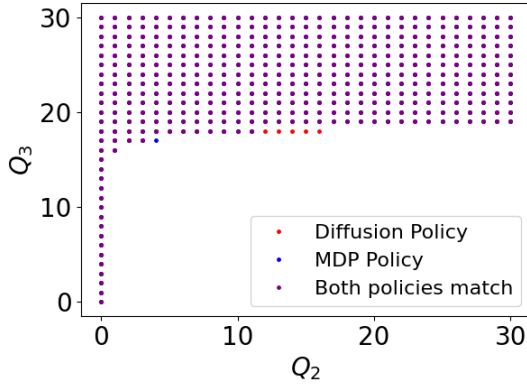
In addition, Table 6 reports the simulated performance (that is, the average present value of all holding costs incurred), with standard errors, of three policies for the original criss-cross example: i)



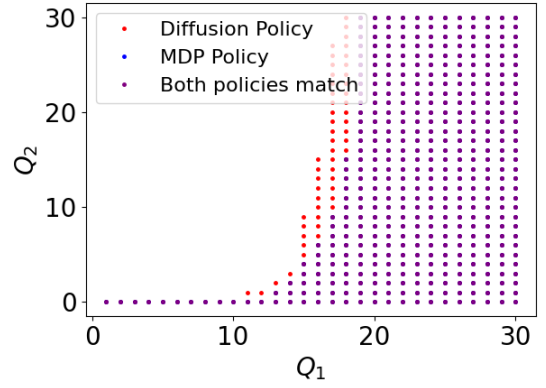
(a) IIB, $Q_1 = 0$



(b) IIC, $Q_3 = 0$



(c) IID-1, $Q_1 = 0$



(d) IID-2, $Q_3 = 0$

Figure 6: Graphical representation of two policies for the criss-cross example: the optimal MDP policy versus our proposed policy based on the singular control approximation.

the policy that gives static priority to class 1 everywhere in the state space, ii) the diffusion policy derived via solution of the approximating singular control problem, and iii) the MDP optimal policy. The performance of the diffusion policy is close to that of the optimal MDP policy (within a fraction of one percent) and much better than that of the static priority policy.

7 Several variants of a three-station example

Harrison [1996] introduced the three-station queueing network displayed in Figure 7, which was further studied by Harrison and Van Mieghem [1997]. It has three external input processes for jobs of types A, B and C. Jobs of types B and C follow the predetermined routes shown in Figure 7. For jobs of type A, there are two processing modes available, and the routing decision must be made irrevocably upon arrival. If a type A job is directed to the lower route, it requires just a single service at station 3, but if it is directed to the upper route, then two services are required at stations 1 and

2, as shown in Figure 7.

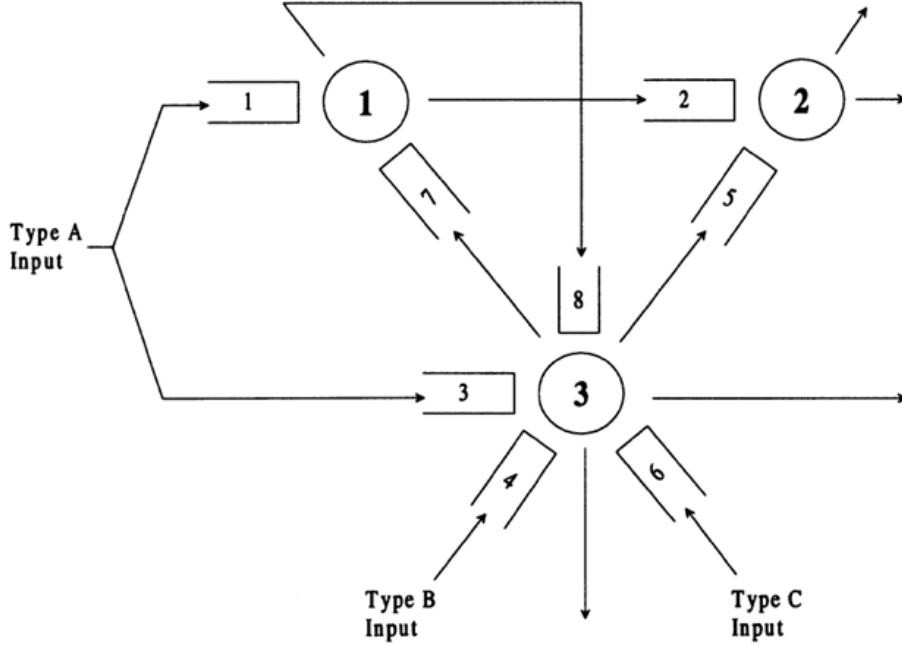


Figure 7: Picture of the 3-station queueing network (copied from Harrison [1996]).

We define a different job class for each combination of job type and stage of completion. Jobs of each class wait for service in a designated buffer, represented by the open-ended rectangles in Figure 7. There are eight classes in total, and each station serves multiple classes. Thus, the system manager must make dynamic sequencing decisions. She also makes dynamic routing decisions to determine which route to use for each type A job. In addition to the dynamic routing and sequencing decisions, the system manager can reject new jobs of any type upon arrival. On the one hand, she incurs a penalty for each rejected job, and on the other hand, there are linear holding costs that provide a motivation for rejecting new arrivals when queues are large. We denote by h_1, \dots, h_8 the holding cost rates (expressed in units like dollars per hour) for the eight job classes identified in Figure 7.

Jobs of types A, B, and C arrive to the system according to independent Poisson processes with rates 0.5, 0.25 and 0.25 jobs per hour, respectively. In addition, class k jobs have exponentially distributed service times with mean m_k where

$$(m_1, \dots, m_8) = (2, 3, 1, 1, 1, 1, 2, 1).$$

As an initial routing mechanism, to be modified by the system manager's dynamic control actions, we suppose that routing decisions for type A arrivals are made by independent coin flips, resulting in the following vector of average external arrival rates into the various classes:

$$(\lambda_1, \dots, \lambda_8) = \left(\frac{1}{4}, 0, \frac{1}{4}, \frac{1}{4}, 0, \frac{1}{4}, 0, 0 \right).$$

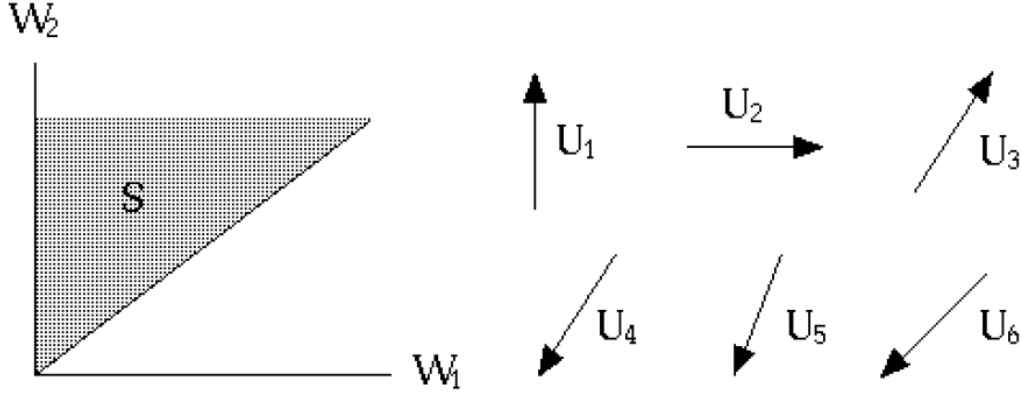


Figure 8: The set S and directions of available controls (copied from Harrison [1996]).

Harrison [1996] and Harrison and Van Mieghem [1997] derived the approximating Brownian control problem and its equivalent workload formulation (EWF) for the three-station queueing network. The EWF will be recapped in this section and then solved numerically for several different cost structures. It is a singular control problem *without* exogenous reflection at the boundary, defined via (9)- (13) in Section 2, except that its state space is not an orthant. Rather, it is the wedge

$$S = \{w \in \mathbb{R}^2 : w = Mz, z \geq 0\}$$

pictured in Figure 8, where M is the 2×8 workload profile matrix

$$M = \begin{bmatrix} 2 & 0 & 2 & 2 & 0 & 6 & 4 & 2 \\ 3 & 3 & 3 & 4 & 1 & 6 & 3 & 3 \end{bmatrix}.$$

As in Section 6, the workload process W for our three-station example is defined by the relationship $W(t) = MZ(t)$, $t \geq 0$, where Z is a scaled queue length process with one component for each job class or buffer, so Z is eight dimensional for the current example. In Section 6 we saw that the workload dimension of the criss-cross network (that is, the dimension of its EWF) equals the number of that system's service stations, but in the current example the system manager's dynamic routing capability reduces the workload dimension from three (the number of service stations) to two. Harrison and Van Mieghem [1997] interpret components of W as scaled workloads for two overlapping "server pools," one composed of servers 1 and 3, the other composed of servers 2 and 3.

In addition to the state space S , Figure 8 shows the six directions of control that are available by increasing different components of the control process U . Those six directions of control are the columns of the matrix

$$G = \begin{bmatrix} 1 & 0 & 2 & -1 & -1/2 & -3/2 \\ 0 & 1 & 3 & -3/2 & -1 & -3/2 \end{bmatrix}.$$

Control modes 1 through 3 correspond to idling servers 1 through 3, while control modes 4, 5 and 6

correspond to rejecting arrivals of types A, B and C, respectively. The associated control cost vector c has the form $c = (0, 0, 0, c_4, c_5, c_6)$, where c_4, c_5 and c_6 will be given strictly positive values in the numerical examples to follow. The three zero components of c reflect an assumption that there is no direct cost of idling servers, whereas the positive values of c_4, c_5 and c_6 represent the costs (expressed in units like dollars per job) to reject new arrivals of types A, B and C, respectively; one can think of those "costs" either as direct penalties or as opportunity costs for business foregone. It is noteworthy that, because the third column of G can be written as a positive linear combination of the first and second columns, and furthermore, control modes 1, 2 and 3 are all costless (that is, $c_1 = c_2 = c_3 = 0$), control mode 3 can actually be eliminated from the problem formulation. Equivalently, one can set $U_3(t) = 0$ for all $t \geq 0$ without loss of optimality, which is consistent with the optimal policies derived below.

The final datum for the exact version of our three-station network control problem is the following vector of holding cost parameters for the eight job classes:

$$h = (h_1, \dots, h_8) = (6, 3, 6, 6, 1, 12, 7, 6).$$

We now define the holding cost function $h(\cdot)$ for our EWF as in Section 6, that is,

$$h(w) = \min\{h \cdot z : Mz = w, z \geq 0\}, \quad w \in S, \quad (33)$$

and with the holding cost parameters specified above, it can be verified that (7) is equivalent to

$$h(w) = w_1 + w_2, \quad w \in S. \quad (34)$$

Let us now consider the singular control problem that is the EWF for our three-station example. Its state space S , control matrix G , holding cost function $h(w)$, and control cost vector c are as specified above. Next, the underlying Brownian motion $X = \{X(t), t \geq 0\}$ that appears in equations (9)-(10) is two-dimensional, with drift vector $\xi = 0$ and covariance matrix

$$A = \begin{bmatrix} 50 & 54 \\ 54 & 69 \end{bmatrix}.$$

Finally, we take the interest rate for discounting in the EWF to be $\gamma = 0.1$.

To formulate or specify the drift control problem by which we approximate the EWF (see Section 3), observe that the first two columns of the control matrix G form the 2×2 submatrix

$$R = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix},$$

which satisfies assumption (12) of our singular control formulation. To accommodate the non-orthant state space of our current example, we must take care to specify which column of R provides the

”back-up reflection capability” (see Section 3 for an explanation of that phrase) at each of the rays that form the boundary of S . For that purpose, we associate the first column of R with the vertical axis in Figure 8 and associate its second column with the oblique ray that forms the other boundary of S . It is crucial, of course, that each column of R points into the interior of S from the boundary ray with which it is associated. In fact, it can be verified that each of the two columns of R is the *unique* column in G that points into the interior of S from its corresponding boundary ray. (Recall that we set $U_3 \equiv 0$ without loss of optimality.)

After making minor changes in the code that implements our computational method, so that it treats a problem in the wedge rather than the quadrant, and setting the upper bound on allowable drift rates to $b = 200$, we solved the singular control problem for the three combinations of cost parameters specified in Table 7. In addition, the same three singular control problems were solved

	c_4	c_5	c_6
Case 1	2	1	1
Case 2	2	1	1
Case 3	1.65	1	2.25

Table 7: The cost parameters.

using the method proposed by Kushner and Martins [1991]. See Figures 9, 10, and 11 for graphical comparisons of the policies derived with the two methods in Cases 1, 2 and 3, respectively.

The solutions produced by the two computational methods are very similar visually, and they share the following gross properties: (i) in all three cases, servers 1 and 2 are idled only when there is no work for them anywhere in the system, and server 3 is never idled; (ii) in Case 1, new arrivals of type C are rejected under certain circumstances, but those of types A and B never are; (iii) in Case 2, new arrivals of types B and C are rejected in different circumstances, but those of type A never are; and in Case 3, new arrivals of all three types are rejected in different circumstances.

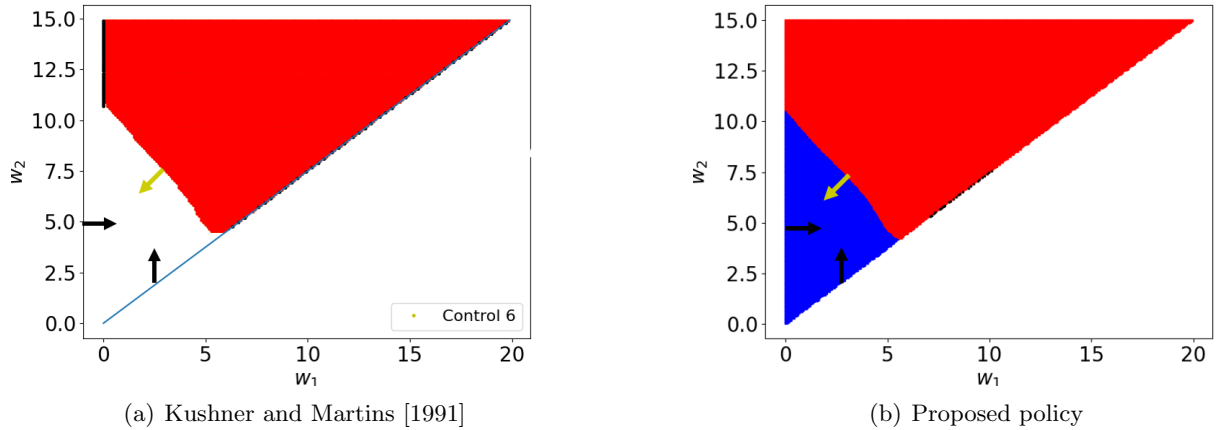


Figure 9: Graphical representation of the policies in case 1.

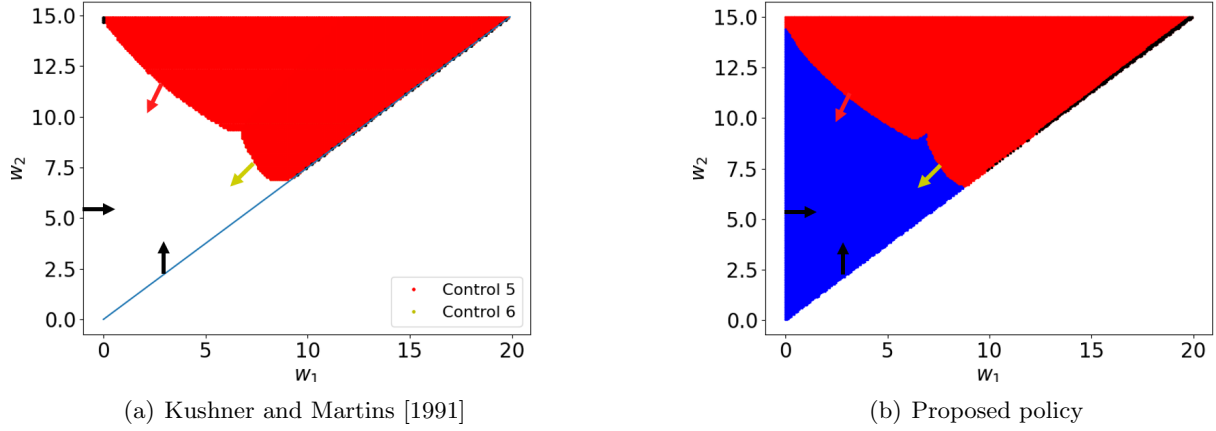


Figure 10: Graphical representation of the policies in case 2.

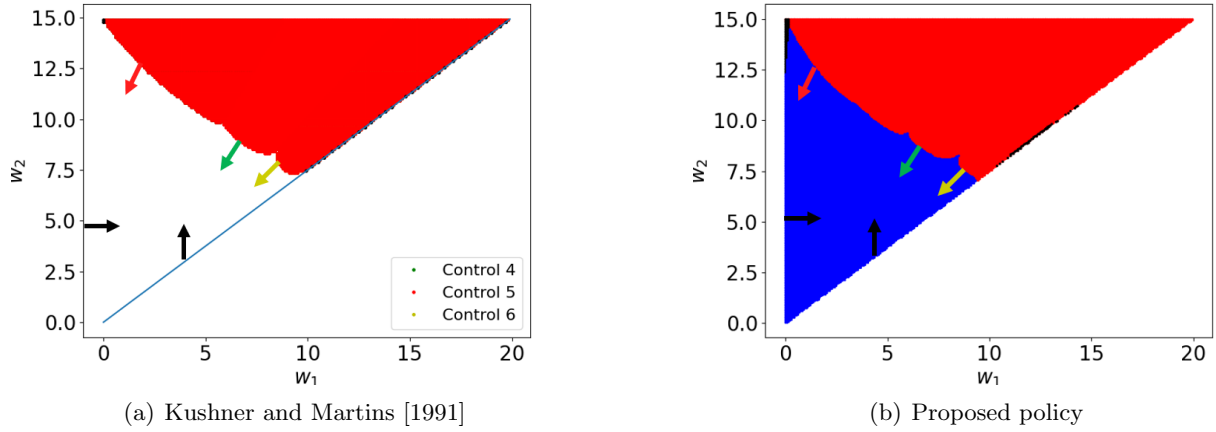


Figure 11: Graphical representation of the policies in case 3.

Finally, Table 8 reports the simulated performance (that is, the average present value of all costs incurred) in all three cases for policies derived via the two methods, assuming both buffers are initially empty, along with the associated standard errors. For comparison, Table 8 also shows the performance of the policy that never turns away jobs and never idles any server except when such idleness is unavoidable. It should be emphasized that these performance estimates were derived by simulating sample paths *of the Brownian singular control problem itself*, not by simulating performance in our original queueing network, as was the case in Sections 5 and 6. The reason for this is that, for our three-station network example, it is not obvious how to interpret or implement the EWF solution in our original problem context, so we postpone that task to future research. See Section 9 for further discussion.

	No control	Proposed policy	Kushner and Martins [1991]
Case 1	102.7 ± 0.12	67.7 ± 0.07	67.7 ± 0.07
Case 2	102.7 ± 0.12	84.2 ± 0.10	84.2 ± 0.10
Case 3	102.7 ± 0.12	85.2 ± 0.07	85.1 ± 0.07

Table 8: Simulation performance for three class of a three-station example.

8 Many queues in series

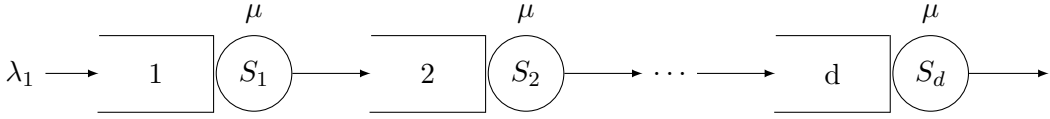


Figure 12: A network of tandem queues.

Figure 12 pictures a network model with d job classes and d single-server stations, extending the tandem queues example presented in Section 5. Class 1 jobs arrive from outside the system according to a Poisson process with rate λ , and they are served at station 1 on a FIFO basis. After completing service at the i^{th} station, class i jobs become class $i + 1$ jobs and proceed to station $i + 1$, where they are again served on a FIFO basis, for $i = 1, \dots, d - 1$. Class d jobs leave the system after completing service. Service times at each station are i.i.d. exponential random variables with mean $m < \lambda^{-1}$, and any service can be interrupted at any time and resumed later without any efficiency loss.

As in Section 5, we assume that any server can be idled at any time, and there is no other mode of control available. This leads to a d -dimensional EWF of the form (9)-(13) with the parameter values specified below. Adopting the sequence index $n = 400$ gives the drift vector

$$\xi = (\sqrt{n}(\lambda - m^{-1}), 0, \dots, 0) = (-1, 0, \dots, 0).$$

The covariance matrix is

$$A = \begin{bmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & \ddots & \ddots & \\ & & \ddots & \ddots & -1 \\ & & & -1 & 2 \end{bmatrix},$$

and the control matrix is

$$G = \begin{bmatrix} 1 & & & & \\ -1 & 1 & & & \\ & -1 & \ddots & & \\ & & \ddots & \ddots & \\ & & & -1 & 1 \end{bmatrix}.$$

As in Section 5, we assume the holding cost function is linear, that is,

$$h(w) = h_1 w_1 + \cdots + h_d w_d, \quad w \in \mathbb{R}_+^d.$$

The numerical values for holding cost rates h_1, \dots, h_d will be specified below. The control cost vector is $c = 0$ (that is, there is no direct cost to idle any server), and the interest rate for discounting in the original model is $r = 0.01$, so the interest rate for our EWF is

$$\gamma = nr = 4.$$

Using our computational method, one can solve problems up to dimension $d \geq 30$, but solving the exact MDP formulation for such problems is not feasible computationally. Therefore, we look for effective benchmark policies, but tuning benchmarks is often computationally demanding as well. Thus, to ease the search for an effective benchmark, we assume d is even and make the following assumption on the holding cost rate parameters:

$$h_2 > h_1 > h_4 > h_3 > \cdots > h_{2k} > h_{2k-1} > \cdots > h_d > h_{d-1}. \quad (35)$$

It is intuitively clear that if $h_j > h_l$ for all $l > j$ and $j = 1, \dots, d-1$, then the control for station j should be minimal, that is, server j should keep working unless its buffer is empty. Therefore, we restrict attention to benchmark policies under which the control for even numbered stations is minimal, that is, servers $2, 4, \dots, d$ are idled only when their buffers are empty. We further restrict attention to benchmark policies under which the control for an odd-numbered station depends only on its own workload and the workload of the following station. To be more specific, for an odd numbered station station i , its server idles at time $t \geq 0$ if either its buffer is empty or the following holds:

$$\beta_i W_i(t) + 1 \leq \beta_{i+1} W_{i+1}(t).$$

We refer to this benchmark policy as a linear boundary policy (LBP) and note that it is similar to the benchmark policy considered in Section 6.4 of Ata et al. [2023].

For our numerical example, we let $d = 6$ and

$$h = (3, 3.9, 2, 2.9, 1, 1.9).$$

Applying our computational method to this singular control problem, we set the upper bound on the feasible drift rates at $b = 20$. The resulting optimal policy is illustrated in Figure 13 and compared with the best linear boundary policy. More precisely, the comparison is for states $(w_1, w_2, 1, 1, 1, 1)$, $(1, 1, w_3, w_4, 1, 1)$ and $(1, 1, 1, 1, w_5, w_6)$ in panels (a), (b) and (c) of Figure 13, respectively. The policy derived from our method closely aligns with the best linear boundary policy. It takes about ten hours of computing to find the optimal policy in our method using a 30-CPU core

machine.

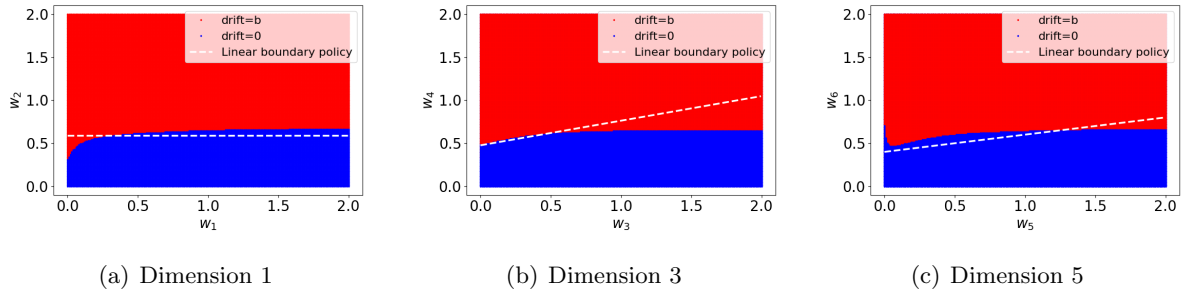


Figure 13: Optimal linear boundary policy and our proposed policy for six queues in series.

Table 9 presents the simulated performance (that is, the average present value of holding costs incurred over the infinite planning horizon) for both policies implemented in our original queueing network, including standard errors for comparison; see Appendix C for a precise description of our proposed policy. Additionally, Table 9 shows the performance of the policy that never idles any servers unless its buffer is empty. These simulations show that the policy we propose, based on the singular control approximation, performs about as well as the best linear boundary policy.

Never idle	Linear boundary policy	Proposed policy
7011 ± 2.8	6924 ± 2.7	6933 ± 2.7

Table 9: Simulation performance for six queues in series.

9 Concluding remarks

Expanding on a statement made in Section 1 of this paper, one may say that there are three steps involved in formulating a queueing network control policy based on a heavy traffic singular control approximation: (a) formulate the approximating singular control problem and derive its equivalent workload formulation (EWF), specifying the data of the EWF in terms of original system data; (b) solve the EWF numerically; and (c) translate the EWF solution into an implementable policy for the original discrete-flow model. The first of those steps has been the subject of research over the past 35 years, starting with Harrison [1988] and including Harrison and Van Mieghem [1997]. The second step is the subject of this paper. For three of the four queueing network examples treated in this paper, the third step (or at least one version of the third step) is spelled out in Appendix C below, but we have not presented a general resolution of the translation problem.

Harrison [1996] proposes a general framework for the translation task using discrete-review policies. To be specific, the proposed policies review the system status at discrete points in time and solve a linear program to make control decisions for the original queueing network. In addition to the first-order problem data, the linear program also uses the gradient ∇V of the optimal value func-

tion for the approximating EWF. Harrison [1996] does not attempt a proof of asymptotic optimality of the proposed policy, but in a follow up paper (Harrison [1998]), the author considers a simple parallel-server queueing network and shows that the framework proposed in Harrison [1996] yields a policy for the original queueing network that is asymptotically optimal in the heavy traffic regime; also see Harrison and Lopez [1999]. For the analysis of discrete-review policies, the state-of-the-art is Ata and Kumar [2005]. The authors prove the asymptotic optimality of a discrete-review policy for a general queueing network under the so called complete resource pooling assumption.

A common feature of the formulations in these papers is that their approximating EWF is one-dimensional, and it admits a pathwise optimal solution. In particular, one can drop the gradient ∇V of its value function from the aforementioned linear program that is the crux of the framework proposed in Harrison [1996]. On the other hand, the solution approach proposed in this paper broadens the class of singular control problems one can solve significantly. As such, it paves the way for establishing the asymptotic optimality of the discrete-review policies in full generality as proposed in Harrison [1996], providing a general resolution of the translation problem that we intend to revisit in future work.

All of the queueing network examples treated in this paper assume independent Poisson arrival streams and exponential service time distributions, but a major virtue of heavy traffic diffusion approximations is that they are distribution free. That is, the approximating diffusion model (in our setting, the approximating singular control problem) depends on the distributions of the original queueing network model only through their first two moments. Our primary reason for using Poisson arrivals and exponential service time distributions in our examples is to allow exact MDP formulations of the associated queueing control problems, which can then be solved numerically for comparison against our diffusion approximations.

The following two paragraphs describe extensions of the problem formulation propounded in Section 2 of this paper, both of which are of interest for queueing applications. Each of them is straight-forward in principle but would require substantial effort to modify existing code.

Weaker assumptions on the structure of the control matrix G . To ensure the existence of feasible controls, we assumed in Section 2 that the first d columns of the control matrix G constitute a Minkowski matrix. That is a relatively strong assumption carried over from [Ata et al., 2023], whose computational method we use to solve our approximating drift control problem. A weaker assumption is that those first d columns constitute a *completely-S* matrix, which Taylor and Williams [1993] showed is a sufficient condition for existence of feasible controls in our setting. It seems very likely that our computational method for drift control, and hence our problem formulation in Section 2, can be generalized to the completely-S case, and that extension is definitely of interest for queueing applications.

More general state space for the controlled process W . To accommodate closed network models, or networks whose total population is bounded, as well as finite buffers and other queueing model features, one would like to allow a state space for W that is a general d -dimensional polytope. This extension also appears to be routine in principle, subject to the following two restrictions: first,

the polytope is *simple* in the sense that no more than d boundary hyperplanes intersect at any of its vertices; and second, at every point on the boundary of the polytope, there exists a convex combination of available directions of control that points into the interior from that point. The latter requirement can be re-expressed in terms of completely-S matrices.

Lastly, we focused only on the discounted cost criterion for brevity, but our approach can be used to solve the singular control problem in the ergodic case too, using the corresponding algorithm of Ata et al. [2023] and making minor changes to our overall approach and the code.

References

- B. Ata and S Kumar. Heavy traffic analysis of open processing networks with complete resource pooling: Asymptotic optimality of discrete review policies. *The Annals of Applied Probability*, 15(1A):331–391, 2005.
- Baris Ata, J Michael Harrison, and Nian Si. Drift control of high-dimensional rbm: A computational method based on neural networks. *arXiv preprint arXiv:2309.11651*, 2023.
- Václav E Beneš, Larry A Shepp, and Hans S Witsenhausen. Some solvable stochastic control problemst. *Stochastics: An International Journal of Probability and Stochastic Processes*, 4(1):39–83, 1980.
- Amarjit Budhiraja and Arka Prasanna Ghosh. A large deviations approach to asymptotically optimal control of crisscross network in heavy traffic. 2005.
- Amarjit Budhiraja and Kevin Ross. Optimal stopping and free boundary characterizations for some brownian control problems. 2008.
- Amarjit Budhiraja, Xin Liu, and Subhamay Saha. Construction of asymptotically optimal control for crisscross network from a free boundary problem. *Stochastic Systems*, 6(2):459–518, 2017.
- Hong Chen, Ping Yang, and David D Yao. Control and scheduling in a two-station queueing network: Optimal policies and heuristics. *Queueing Systems*, 18:301–332, 1994.
- JG Dai and J Michael Harrison. *Processing networks: fluid models and stability*. Cambridge University Press, 2020.
- J Michael Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications*, pages 147–186. Springer, 1988.
- J Michael Harrison. The bigstep approach to flow management in stochastic processing networks. *Stochastic Networks: Theory and Applications*, 4:147–186, 1996.
- J Michael Harrison. Heavy traffic analysis of a system with parallel servers: asymptotic optimality of discrete-review policies. *The Annals of Applied Probability*, 8(3):822–848, 1998.

- J Michael Harrison and Martin I Reiman. Reflected brownian motion on an orthant. *The Annals of Probability*, 9(2):302–308, 1981.
- J Michael Harrison and Michael I Taksar. Instantaneous control of brownian motion. *Mathematics of Operations research*, 8(3):439–453, 1983.
- J Michael Harrison and Jan A Van Mieghem. Dynamic control of brownian networks: state space collapse and equivalent workload formulations. *The Annals of Applied Probability*, pages 747–771, 1997.
- J Michael Harrison and Lawrence M Wein. Scheduling networks of queues: heavy traffic analysis of a simple open network. *Queueing Systems*, 5:265–279, 1989.
- J.M. Harrison and M.J. Lopez. Heavy traffic resource pooling in parallel-server systems. *Queueing Systems*, 33:339–368, 1999.
- Ioannis Karatzas. A class of singular stochastic control problems. *Advances in Applied Probability*, 15(2):225–254, 1983.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Sunil Kumar and Kumar Muthuraman. A numerical method for solving singular stochastic control problems. *Operations Research*, 52(4):563–582, 2004.
- Harold J Kushner. Numerical methods for stochastic control problems in continuous time. *SIAM Journal on Control and Optimization*, 28(5):999–1048, 1990.
- Harold J Kushner and Luiz Felipe Martins. Numerical methods for stochastic singular control problems. *SIAM journal on control and optimization*, 29(6):1443–1475, 1991.
- Luis Filipe Martins, Steven E Shreve, and H Mete Soner. Heavy traffic convergence of a controlled, multiclass queueing system. *SIAM journal on control and optimization*, 34(6):2133–2171, 1996.
- Andrinandrasana David Rasamoelina, Fouzia Adjailia, and Peter Sinčák. A review of activation function for artificial neural network. In *2020 IEEE 18th World Symposium on Applied Machine Intelligence and Informatics (SAMI)*, pages 281–286. IEEE, 2020.
- Lisa Maria Taylor and Ruth J Williams. Existence and uniqueness of semimartingale reflecting brownian motions in an orthant. *Probability Theory and Related Fields*, 96(3):283–317, 1993.

A Implementation details

A.1 Our policies

We first add details about the algorithm used to solve the approximating drift control problems described in Section 3. That algorithm, which was the subject of our earlier paper Ata et al. [2023], makes heavy use of neural networks. The problem-specific hyper-parameters that were used in training the neural networks for our examples are listed in Table 10, and the hyper-parameters that are common to all our examples are specified immediately below.

Common hyperparameters. We let the batch size $B = 256$, the time horizon $T = 0.1$, and the discretization step size $\Delta t = 0.1/64$.

Optimizer. We used the Adam optimizer [Kingma and Ba, 2014].

Activation function. We use the 'elu' action function [Rasamoelina et al., 2020].

We modified the algorithm described in Ata et al. [2023] to speed up the neural network training. These modifications exploit the special structure of our examples and are described below.

Hyperparameters	Parallel	Tandem	Criss-cross	Three-station	Many queues
Section	4	5	6	7	8
Dimension	30	2	2	2	6
b	10	20	20	200	20
Reference policy μ	-1.0	-1.0	-0.5	-5.0	-0.5
#Iterations	6000	6000	6000	6000	9000
#Epochs	135	34	34	40	52
	5e-4 (0,9.5e3)	5e-4 (0,3e3)	5e-4 (0,3e3)	5e-4 (0,1.9e4)	5e-4 (0,9.5e3)
	3e-4 (9.5e3,2.2e4)	3e-4 (3e3,9e3)	3e-4 (3e3,9e3)	3e-4 (1.9e4,4.4e4)	3e-4 (3e3,9e3)
Learning rate scheme	1e-4 (2.2e4, ∞)	1e-4 (9e3, ∞)	1e-4 (9e3, ∞)	1e-4 (4.4e4,7e4)	1e-4 (9e3,1e5)
				3e-5 (7e4, ∞)	5e-5 (1e5,1.5e5)
					5e-2 (1.5e5, ∞)
#Hidden layers	3	3	3	3	3
#Neurons in each layer	300	50	50	100	100

Table 10: Hyperparameters used in the numerical examples.

Decay loss. For the parallel queue example treated in Section 4, we implement the same decay loss approach that was used in Section E.1 of Ata et al. [2023]. Here we let

$$\tilde{b} = \left(7 - \frac{\text{iteration}}{4800}\right)^+.$$

Increasing upper bound b . For the criss-cross network (Section 6), the three-station example (Section 7), and the many queues in series example (Section 8), we gradually increase the upper bound b during the neural network training. We start from a small value and gradually increase it to the target level. To be specific, we let

$$\hat{b} = \begin{cases} \min(20, 4 + \frac{\text{iteration}}{80}) & \text{for the criss-cross network,} \\ \min(200, 40 + \frac{\text{iteration}}{800}) & \text{for the three-station network,} \\ \min(20, 4 + \frac{\text{iteration}}{14400}) & \text{for the many queues in series example.} \end{cases}$$

Shape constraints. Let $V : \mathbb{R}_+^d \rightarrow \mathbb{R}$ denote the optimal value function for the singular control problem (9)-(13) introduced in Section 2. That is, let $V(w)$ denote the minimum achievable expected present value of costs incurred over the infinite planning horizon, given that the initial state vector is $W(0) = w$. The computational method that we proposed in Ata et al. [2023] provides not only an approximation for V but also an approximation $G : \mathbb{R}_+^d \rightarrow \mathbb{R}^d$ for the gradient of V . In both cases, the approximation comes in the form of a neural network. In order to emphasize this, we will write $G^\theta(w) = (G_1^\theta(w), \dots, G_d^\theta(w))$ where θ denotes the neural network parameters.

For the examples considered in Sections 5- 8, one can infer certain properties of the value function and its gradient, which we use to speed up the training process. Next, we describe these properties, which we refer to as the shape constraints, for each example.

- **Criss-cross network example.** In all four cases considered in Section 6, one observes the following from Table 5: i) When $2w_1 > w_2$, the holding cost function $h(w_1, w_2)$ is increasing in w_1 , and ii) when $2w_1 \leq w_2$, it is increasing in w_2 . Intuitively, these suggest that

$$\begin{aligned} \frac{\partial V(w)}{\partial w_1} &> 0 \quad \text{when } 2w_1 > w_2, \\ \frac{\partial V(w)}{\partial w_2} &> 0 \quad \text{when } 2w_1 \leq w_2. \end{aligned}$$

Thus, we seek a neural network approximation that satisfies these "shape constraints" or monotonicity requirements. That is, we seek a neural network approximation that satisfies

$$\begin{aligned} G_1^\theta(w) &> 0 \quad \text{when } 2w_1 > w_2, \\ G_2^\theta(w) &> 0 \quad \text{when } 2w_1 \leq w_2. \end{aligned}$$

We pursue this goal by amending our loss function for the neural network training with the following penalty term:

$$\left([G_1^\theta(w)]^- \mathbb{I}\{2w_1 > w_2\} + [G_2^\theta(w)]^- \mathbb{I}\{2w_1 \leq w_2\} \right)^2,$$

where $\mathbb{I}\{\cdot\}$ is the indicator function.

- **Three-station network example.** Recall that the holding cost function $h(w) = w_1 + w_2$ for this example. Because $h(w)$ is monotone, one expects intuitively that $\nabla V(w) \geq 0$. Thus, we look for a neural network approximation that has $G^\theta(w) \geq 0$ for $w \in S$. To do so, we add the

following penalty term to the loss function used for neural network training:

$$\left([G_1^\theta(w)]^-\right)^2 + \left([G_2^\theta(w)]^-\right)^2.$$

- **Tandem queue and many queues in series examples.** As described in Appendix C, server i ($i = 1, \dots, d-1$) idles if

$$\frac{\partial V(w)}{\partial w_i} - \frac{\partial V(w)}{\partial w_{i+1}} < 0. \quad (36)$$

Intuitively, idling server i becomes less appealing if its workload increases. Thus, one expects the left-hand side of (36) to be increasing in w_i . That is,

$$\frac{\partial}{\partial w_i} \left[\frac{\partial V(w)}{\partial w_i} - \frac{\partial V(w)}{\partial w_{i+1}} \right] > 0.$$

Therefore, we seek a neural-network approximation that satisfies

$$\frac{\partial}{\partial w_i} [G_i^\theta(w) - G_{i+1}^\theta(w)] > 0, \quad i = 1, \dots, d-1.$$

To do so, we add the following penalty term to the loss function

$$\sum_{i=1}^{d-1} \left[\left(\frac{\partial}{\partial w_i} [G_i^\theta(w) - G_{i+1}^\theta(w)] \right)^- \right]^2.$$

A.2 MDP solutions

For the tandem queues example (Section 5) and the criss-cross network example (Section 6), we solve the MDP associated with the original queueing network models via value iteration. For each example, we truncate the state space by imposing an upper bound for each buffer. To be specific, we let the upper bound be 1000 for the tandem queues example and 300 for the criss-cross example. In both cases, we stop the value iteration when the iteration error is less than $\epsilon = 0.1$.

A.3 A benchmark for the three-station example via Kushner and Martins [1991]

In addition to solving the singular control problem for the three-station network example (Section 7) using our method, we also solve it by adopting the method of Kushner and Martins [1991], which builds on Kushner [1990]. This provides a benchmark for comparison with our method. Note that our covariance matrix does not satisfy the assumption (5.5) in Kushner [1990]. Therefore, we used different grid sizes in different dimensions that satisfy (5.9) in Kushner [1990]. Specifically, the discretization grid (using the notation in Kushner [1990]) is defined with $h_1 = 0.1$ in the first dimension and $h_2 = 0.108$ in the second dimension; and the state space is truncated at $400h_1 = 40$ in the first dimension, and it is truncated at $400h_2 = 43.2$ in the second dimension.

A.4 Linear boundary policies

As mentioned in Section 8, we use the linear boundary policy (LBP) as the benchmark for the many queues in series example. To design an effective LBP, we perform a grid search over its parameters β_1, \dots, β_6 . In particular, we consider the following values:

$$\begin{aligned}\beta_i &\in \{0, 0.1, \dots, 1.0\} \text{ for } i = 1, 3, 5, \\ \beta_i &\in \{0.4, 0.5, \dots, 3.6\} \text{ for } i = 2, 4, 6.\end{aligned}$$

The range of values considered are informed by a preliminary search over a broader range of values. A brute-force search over all possible parameter combinations involves evaluating $(11 \times 33)^3$ scenarios, which is prohibitive computationally. Therefore, we proceed with the following heuristic approach:

1. We search β_1 and β_2 assuming servers 3 and 5 are non-idling.
2. Given the best performing β_1 and β_2 obtained in the previous step, we then search for β_3 and β_4 assuming server 5 is non-idling.
3. Given the best performing β_1, \dots, β_4 obtained in the previous two steps, we then search for β_5 and β_6 .

This heuristic method reduces the number of cases to $3 \times 11 \times 33$. For each case, we simulate 300000 independent paths. The resulting parameters are as follows:

$$(\beta_1, \dots, \beta_6) = (0.0, 1.7, 0.6, 2.1, 0.5, 2.5).$$

B Details of the simulation studies for policy comparisons

We performed 400,000 independent replications for all examples considered except for the parallel-server queueing network example (Section 4). For that example, we performed 100,000 replications.

Further details for the tandem queues, the criss-cross network and the many queues in series examples. For these examples, we simulate the original queueing network models, which avoids discretization errors. Also, we choose the time horizon for the simulation sufficiently long to approximate the infinite horizon discounted performance. To be specific, we make sure each simulated path is run longer than 1400 time units. Because the interest rate $r = 0.01$, the discount factor at the termination is less than or equal to $\exp -1400r \approx 8e^{-7}$, indicating a negligible error due to truncation of the time horizon.

Further details for the parallel-server queueing network and the three-station queueing network examples. For those examples, we simulate the approximating diffusion models (Section 2) via discretizing the time of the diffusion model; see Table 11. Then the singular control is implemented using a unit-step control. To be specific, if the process crosses the state boundary, we repeatedly apply a unit-step control to push it back into state space until that happens.

	Parallel-server queueing network	Three-station network
Time discretization	3.125e-5	0.0015625
Step size for unit control	0.01	0.1 (1st dim), 0.108 (2nd dim)

Table 11: The specifics of the discretization grids and control unit steps for the parallel queues and three-station queues

C Scheduling policies for queueing network examples.

Recall from Appendix A.1 that $V : \mathbb{R}_+^d \rightarrow \mathbb{R}$ denotes the optimal value function for the singular control problem (9)-(13) introduced in Section 2. That is, $V(w)$ denotes the minimum achievable expected present value of costs incurred over the infinite planning horizon, given that the initial state vector is $W(0) = w$. Also recall that the computational method that we proposed in Ata et al. [2023] provides not only an approximation for V but also an approximation $G : \mathbb{R}_+^d \rightarrow \mathbb{R}^d$ for the gradient of V . (In both cases, the approximation comes in the form of a neural network.) It is this function $G(\cdot) = (G_1(\cdot), \dots, G_d(\cdot))$ that we use in the algorithmic implementation of our proposed scheduling policies for queueing network examples in Sections 5, 6 and 8.

In each case, we observe the system state $Q(t) = q$ immediately after a state transition time $t \geq 0$ (that is, immediately after an arrival from outside the system or a service completion) and determine an action (that is, which servers to idle, if any, and which classes to serve at stations where there is a choice) that remains in place until the next state transition.

To spell out the mechanics of those implementations, we begin with the tandem queues example (Section 5) and many queues in series (Section 8), because their policy descriptions are simplest. In those examples, the last server keeps working unless its buffer is empty. Thus, we will focus on the upstream servers to describe our proposed policy.

Tandem queues example (Section 5). Denoting by w a generic state of the workload process W , the phrase “ w lies in the region colored red in Figure 3” means precisely that $G_1(w) \leq G_2(w)$. In words, this means that an increase in U_1 , which causes a displacement downward and to the right, effects a decrease in $V(\cdot)$, so increasing U_1 is favorable; gradient inequalities are interpreted similarly in the other two examples below. Thus, having observed that $Q(t) = q$ immediately after a state transition time $t \geq 0$, and recalling the definition $W(\cdot) = Q(n\cdot)/\sqrt{n}$, where n is the sequence index or scaling parameter introduced earlier, our computational rule is the following: server 1 keeps working unless either $q_1 = 0$ or $G_1(q/\sqrt{n}) \leq G_2(q/\sqrt{n})$, in which case it idles.

Many queues in series example (Section 8). Suppose that a d -dimensional queue length vector $Q(t) = q$ is observed immediately after a state transition time $t \geq 0$. For $i = 1, \dots, d-1$, server i keeps working unless either $q_i = 0$ or $G_i(q/\sqrt{n}) \leq G_{i+1}(q/\sqrt{n})$, in which case it idles.

Criss-cross network example (Section 6). We use the scheduling policy proposed by Harrison and Wein [1989] and Martins et al. [1996]. As discussed in Section 6, server 2 keeps working as long as its buffer isn’t empty. Focusing attention on server 1, Algorithm 1 describes the scheduling policy. As a preliminary, recall that, given observations of the three-dimensional queue length pro-

cess $Q(\cdot)$, we define $Z(\cdot) = n^{-1/2}Q(n\cdot)$, and define the associated two-dimensional (scaled) workload process $W(\cdot) = MZ(\cdot)$.

Algorithm 1 Scheduling policy for the criss-cross network.

Input: Queue length vector $Q(t) = q$ observed immediately after a state transition time $t \geq 0$, the corresponding workload vector $w = n^{-1/2}Mq$, a safety stock parameter s , and the gradient function $G = (G_1, G_2)$

Output: A control action for server 1

```

1: if  $q_1 = 0$  and  $G_1(w) < 0$  then
2:   idle server 1
3: else
4:   if  $q_3 > s$  or  $G_2(w) < 0$  then
5:     give priority to class 1
6:   else
7:     give priority to class 2
8:   end if
9: end if

```

D Analytical solution of a one-dimensional singular control problem

We consider the one-dimensional singular control problem stated below. We assume $h > rc > 0$. Otherwise, one can show that it is optimal not to exert any control.

$$\begin{aligned}
& \min_{U(\cdot)} E \left[\int_0^{+\infty} e^{-rt} [hW(t) + cdU(t)] \right] \\
& \text{subject to} \\
& W(t) = w + X(t) + Y(t) - U(t), \quad t \geq 0, \\
& Y, U \text{ are adapted to } X, \\
& Y \text{ is continuous and non-decreasing with } Y(0) = 0, \\
& Y \text{ increases at times } t \geq 0 \text{ when } W(t) = 0, \\
& U \text{ is right-continuous and non-decreasing with } U(0) \geq 0, \\
& W(t) \geq 0, \quad t \geq 0.
\end{aligned}$$

Using a smooth-pasting argument one can solve the corresponding HJB equation by considering the following problem; see for example, Harrison and Taksar [1983]: Find a threshold value $w^* > 0$ and $V : \mathbb{R}_+ \rightarrow \mathbb{R}$ that is twice continuously differentiable such that they jointly satisfy

$$\begin{aligned}
& \frac{a}{2} V''(w) + hw = rV(w), \quad \text{for } w \in [0, w^*] \\
& V'(0) = 0 \text{ and } V'(w) = c \text{ for } w \geq w^*.
\end{aligned}$$

One can verify that the solution is of the following form:

$$V^*(w) = \begin{cases} V_1(w) & \text{if } w < w^* \\ V_2(w) & \text{if } w \geq w^* \end{cases},$$

where

$$V_1(w) = \frac{h\sqrt{a}e^{-\frac{\sqrt{2}\sqrt{r}w}{\sqrt{a}}}}{\sqrt{2}r^{3/2}} + \frac{hw}{r} + Ce^{\frac{\sqrt{2}\sqrt{r}w}{\sqrt{a}}} + Ce^{-\frac{\sqrt{2}\sqrt{r}w}{\sqrt{a}}}, \text{ and}$$

$$V_2(w) = cw - cw^* + V_1(w^*),$$

for threshold w^* and constant C we specify next. Furthermore, the optimal control is that $U(t)$ increases only when $W(t) = w^*$ and confines the process in the interval $[0, w^*]$.

To be specific, one needs to solve the following two equations to determine w^* and C :

$$V_1'(w^*) = c \text{ and } V_1''(w^*) = 0.$$

For the parameter values we set forth in Section 4 (that is, $h = 2, a = c = 1, r = 0.1$), we have $w^* = 0.722331$ and $C = -15.3786$.