Queueing Systems 44, 313–363, 2003

In Search of Sensitivity in Network Optimization *

MIKE CHEN, CHARUHAS PANDIT and SEAN MEYN
mikechen@uiuc.edu, cpandit@students.uiuc.edu, s-meyn@uiuc.edu
Department of Electrical and Computer Engineering and the Coordinated Sciences Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Received 4 January 2002; Revised 7 February 2003

Abstract. This paper concerns the following questions regarding policy synthesis in large queueing networks: (i) It is well known that an understanding of variability is important in the determination of safety stocks to prevent unwanted idleness. Is this the only value of high-order statistical information in policy synthesis? (ii) Will a translation of an optimal policy for the deterministic fluid model (in which there is no variability) lead to an allocation which is approximately optimal for the stochastic network? (iii) What are the sources of highest sensitivity in network control? A sensitivity analysis of an associated fluid-model optimal control problem provides an exact dichotomy in (ii). If an optimal policy for the fluid model is 'maximally non-idling', then variability plays a small role in control design. If this condition does not hold, then the 'gap' between the fluid and stochastic optimal policies is exactly proportional to system variability. Furthermore, under mild assumptions, we find that the optimal policy for the stochastic model is closely approximated by an affine shift of the fluid optimal solution. However, sensitivity of steady-state performance with respect to perturbations in the policy vanishes with increasing variability.

Keywords: queueing networks, routing, scheduling, optimal control

AMS subject classification: 90B35, 68M20, 90B15, 93E20

1. Introduction

Control synthesis for network models is of great interest in both academia and industry for obvious reasons. It is equally obvious that direct dynamic-programming approaches to optimization lead to intractable optimality equations for all but the simplest network models. This has led to the development of various alternative network models tuned to address particular issues such as optimal control in 'heavy-traffic'; the impact of breakdowns; or steady-state performance (e.g., [3,11,13,19,20,32,49]).

The simplest model is the linear, deterministic fluid model used, for example, in [2,9,10,18,34–36,45,50,52]. It provides a framework for policy synthesis for large networks based on linear programming methods, and this leads to attractive approaches to sensitivity analysis through associated Lagrange-multiplier techniques.

* This paper is based upon work supported by the National Science Foundation under Award Nos. ECS 99 72957 and DMI 00 85165. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Further motivation for the deterministic network model comes from an emerging theory establishing solidarity among various models. A strong solidarity between stability of fluid models and their stochastic counterparts is established in [14,15]. Results establishing solidarity among respective optimal control solutions are developed in [40,41], based on this stability theory. These results concern the following optimality criteria:

$$J(t) = \int_0^\infty c\bigl(q(t;x)\bigr)\,\mathrm{d}t, \tag{1}$$

$$K(x) = \limsup_{T\to\infty} \frac{1}{T}\int_0^T \mathsf{E}\bigl[c\bigl(Q(t;x)\bigr)\bigr]\,\mathrm{d}t, \tag{2}$$

where $q(t;x)$, $Q(t;x)$ denote the vector of buffer lengths at time $t \geqslant 0$, for the fluid and stochastic models, respectively, with initial conditions $q(Q;x) = Q(0;x) = x$. A policy is called *optimal* if it minimizes the associated value for each $x \in \mathsf{X}$.

It is shown in [40,41] that scaled optimal solutions for the stochastic network are approximated by the optimal solution for the fluid model. Conversely, a policy based on an optimal solution for the fluid model will be approximately optimal for the stochastic network model, provided a certain *effective cost* is monotone [43, theorems 4.3–4.5].

The deterministic fluid model can be refined by the addition of an additive disturbance. When the disturbance is Gaussian then one obtains the Brownian model developed in, for example, [11,20,21,24,29,32,33,37,46,47].

Certain small Brownian network models have yielded to exact analysis, and a translation of the optimal policy to a network model with general statistics is then shown to be approximately optimal by comparison with the Brownian network. A now standard approach to policy translation is to impose thresholds, or safety-stocks. This ensures feasibility of solutions by preventing ‘deadlocks’ or ‘starvation of resources’. One example is the ‘criss-cross network’ introduced in [24], and further studied in several subsequent references. When the effective cost is monotone then one obtains a policy that is approximately pathwise optimal in heavy traffic [4,38].

When the effective cost for the fluid-model is not monotone then an optimal policy is *not* pathwise optimal for the Brownian model (see [43, section 4.5]). An optimal policy for the Brownian model, when it exists, is defined by nonlinear switching-curves in workload-space. In this case only qualitative structural results have been established in small examples (e.g., [38]), and numerical studies have appeared in [17,31].

The aforementioned optimality theory is based upon a workload representation of the network under study. Related general constructions are described in [8,22,43], and this framework will form the basis of the results of reported here.

We highlight here some of the specific issues addressed in the present paper:

(i) How does a policy for a stochastic model change with increasing variability? When the monotone assumption of [43] does not hold we demonstrate in proposition 5.2 that optimal policies scale linearly with variability.

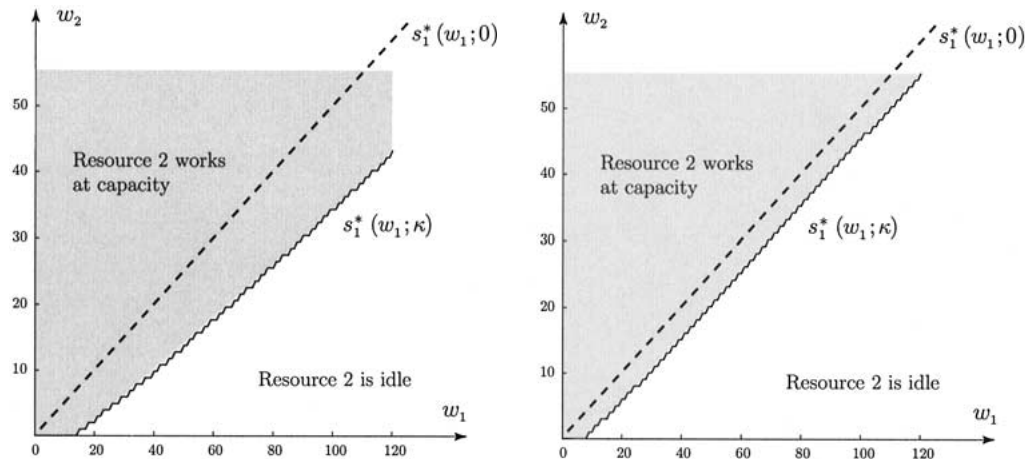

(a) Height process is null recurrent (b) Switching curve is attracting

Figure 1. Optimal switching-curves for the Poisson workload-relaxation in case II obtained using value iteration (see section 6.1). On the left is case II(a), where the network is balanced, and on the right is case II(b). The difference between the fluid and stochastic switching-curves is significant in case II(a). This corresponds to null-recurrence of the associated height process (defined in section 4.1). The optimal policy in case II(b) is accurately approximated by an affine policy whose offset is determined by an associated controlled Brownian model (see (50)).

(ii) When the effective cost is not monotone we show that the discounted-optimal policy for the stochastic model may be approximated by a computable affine translation of the myopic policy for the fluid model (see theorem 4.4). Under certain geometric conditions, the affine policy is pathwise optimal for the fluid model, and under these conditions it is shown that the average-cost optimal policy is also approximated by a computable affine policy (see theorems 4.6 and 4.7).

(iii) In theorem 5.4 we find that, although the policy changes linearly with increased variability, second order sensitivity vanishes as variability approaches infinity.

(iv) In the process of translating a fluid policy to a stochastic model we identify parameters that have a strong impact on performance. These parameters correspond to hard constraints in the deterministic optimal control problem. Some results are described in section 5.3, and illustrated in section 6.

Figure 1(b) shows results from one numerical example used to illustrate the findings of theorem 4.7. The affine shift parameter obtained from consideration of a controlled Brownian model results in a policy that matches almost exactly the optimal policy for a network with Poisson statistics.

The remainder of the paper is organized as follows. Section 2 contains a description of the models used for analysis and control synthesis, and includes a construction of their workload-relaxations. Background on optimal control for these models is provided in section 3. Section 4 develops affine approximations, as described in (ii) above; section 5 contains an investigation of the sensitivity of optimal control solutions with respect to various network parameters, and section 6 provides detailed numerical examples. Conclusions and suggestions for further research are contained in section 7.

2. Network models and their relaxations

The results of this paper are based on two primary network models: A linear, stochastic network model and its fluid counterpart. In this section we provide properties of these models and associated workload models.

Stochastic network models are the focus of most research in the networks area since they capture a range of behaviors. In this sense, the deterministic model is limited. For example, it is obvious that it has little value for steady-state *prediction* since no variability is included in the model. However, the focus of this paper is on optimal control solutions for these various models, and the relationship between their respective control solutions. Both the reduced complexity and linearity of the deterministic fluid model are tremendous virtues in control design.

We begin with a treatment of this model.

2.1. Linear fluid model

In this paper we restrict to the following generalized *scheduling model* in which there are an equal number of activities and customer classes. This is captured in assumption (iv) below. The motivation of (iv) in this paper is to provide a transparent formulation of *workload* for the fluid model, without resorting to the abstract treatments of [22,23,43].

The linear fluid model. The vector of buffer-levels is defined by the vector equation

$$q(t;x) = x + Bz(t;x) + \alpha t, \quad t \in \mathbb{R}_+, \tag{3}$$

where the *state process* $\boldsymbol{q}$ and *allocation process* $\boldsymbol{z}$ evolve on $\mathbb{R}_+^\ell$, for some integer ℓ. They are subject to the following interpretations and constraints:

(i) $q(t;x) \in \mathsf{X}$ is a vector of buffer-levels of various materials in the network, where $\mathsf{X} \subset \mathbb{R}_+^\ell$ is a polyhedron representing both positivity constraints, and bounds on buffer levels if present.

(ii) $\zeta(t;x) := (\mathrm{d}^+/\mathrm{d}t)z(t;x)$ is a vector of instantaneous processing rates of various activities. It is subject to linear constraints:

$$\zeta(t;x) \in \mathsf{U},\ x \in \mathsf{X},\ t \in \mathbb{R}_+, \quad \mathsf{U} := \{u \in \mathbb{R}_+^\ell\colon Cu \leqslant \mathbf{1}\}, \tag{4}$$

where the *constituency matrix* C is an $\ell_r \times \ell$ matrix with binary entries, and $\mathbf{1}$ denotes a vector of ones. For each $1 \leqslant i \leqslant \ell_r$ we define the set of activities associated with this resource via

$$\mathcal{C}_i = \{j\colon C_{i,j} = 1\} \subset \{1, \ldots, \ell\}.$$

An activity may share resources. That is, the rows of C need not be orthogonal.

(iii) The vector $\alpha \in \mathbb{R}^\ell_+$ represents the rate of exogenous arrivals to the network, and possibly also exogenous demands for materials *from* the network.

(iv) The matrix B is of the form

$$B = -(I - R)M, \tag{5}$$

where $M = \mathrm{diag}(\mu)$, with $\mu \in \mathbb{R}^\ell_+$ the vector of *service rates*. The *routing matrix* R has non-negative entries, and we assume that its spectral radius is strictly less than one.

Throughout much of the paper we take $\mathsf{X} = \mathbb{R}^\ell_+$. This restriction will be relaxed in sections 5.3 and 6 so that we may include possible buffer constraints.

To define workload and related concepts, first note that under (iv) the following inverse exists as a power series

$$(I - R)^{-1} \sum_{k=0}^{\infty} R^k.$$

We define the *workload matrix* by

$$\Xi := -CB^{-1} = CM^{-1}(I - R)^{-1},$$

and we let $\{\xi^i\colon 1 \leqslant i \leqslant \ell_r\} \subset \mathbb{R}^\ell$ denote the rows of the workload matrix. These are called *workload vectors*.

The workload vectors satisfy $\xi^i \in \mathbb{R}^\ell_+$ under the assumptions of (iv) since $(I - R)^{-1}$ has non-negative entries. Note that positivity may fail in more general settings. In particular, for demand driven models, positivity is typically impossible (several examples are given in [12]).

We define the *vector load* by

$$\boldsymbol{\rho} = \Xi\alpha, \tag{6}$$

and the *system load* is defined to be the maximum $\rho = \max_i \rho_i$. This definition is motivated by consideration of an equilibrium model. Suppose that the network is in equilibrium with $q(t) \equiv \boldsymbol{\theta}$, where $\boldsymbol{\theta} \in \mathbb{R}^\ell$ is the vector of zeros. It follows that $(\mathrm{d}/\mathrm{d}t)q(t) \equiv \boldsymbol{\theta}$, and consequently there exists at least one solution $\zeta^{\mathrm{ss}} \in \mathsf{U}$ to the equilibrium equation

$$B\zeta^{\mathrm{ss}} = -\alpha. \tag{7}$$

Under the assumptions of (iv) the matrix B is invertible, and hence

$$\zeta^{\mathrm{ss}} = -B^{-1}\alpha.$$

We conclude that $\zeta^{\mathrm{ss}} \in \mathsf{U}$, so that there exists an equilibrium solution, if and only if $\rho \leqslant 1$.

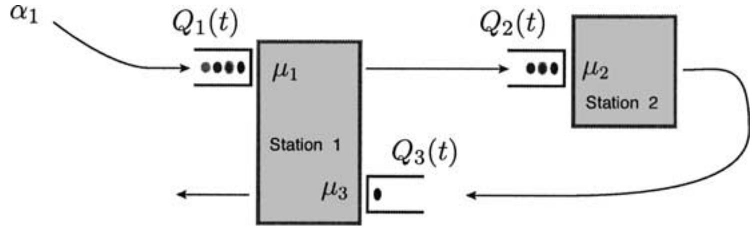

Figure 2. A three-buffer scheduling model.

The workload matrix is also important in describing the dynamic behavior of the network. For this we consider the *workload process*, defined by

$$w(t; w) = \Xi q(t; x), \quad t \geqslant 0, \ w = \Xi x.$$

The following rate constraint on w follows from the definitions.

Proposition 2.1. For any initial condition $x \in \mathsf{X}$, and any feasible trajectory $\boldsymbol{q}$, the corresponding workload process $\boldsymbol{w}$ satisfies the rate constraints,

$$\frac{\mathrm{d}^+}{\mathrm{d}t} w_i(t; i) \geqslant -(1 - \rho_i), \quad 1 \leqslant i \leqslant \ell_r.$$

The three-dimensional example illustrated in figure 2 will be revisited throughout the paper to illustrate our conclusions. The system parameters in this example are

$$B = \begin{bmatrix} -\mu_1 & 0 & 0 \\ \mu_1 & -\mu_2 & 0 \\ 0 & \mu_2 & -\mu_3 \end{bmatrix}, \qquad \alpha = \begin{bmatrix} \alpha_1 \\ 0 \\ 0 \end{bmatrix}, \qquad C = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \tag{8}$$

and the workload vectors and load parameters are given by

$$\begin{aligned} \xi^1 &= (m_1 + m_3, m_3, m_3)^{\mathrm{T}}, \qquad & \rho_1 &= \alpha_1(m_1 + m_3), \\ \xi^2 &= (m_2, m_2, 0)^{\mathrm{T}}, & \rho_2 &= \alpha_1 m_2, \end{aligned}$$

where $m_i = 1/\mu_1$, $i = 1, 2, 3$, and we have used $\boldsymbol{\rho} = \Xi\alpha \in \mathbb{R}^2_+$, with Ξ given in (6).

2.2. *The controlled random-walk model*

We describe here a special case of the linear stochastic model introduced in [28]. Our objective is to obtain an appropriate *Markov model* for the physical system of interest. A common approximation technique is the use of Erlang distributions to obtain a model of this form. While this approach does give a Markov model, this is at the expense of an unnaturally complex state space.

To motivate the controlled random-walk model we begin with a description of the M/M/1 queue. It is well known that through appropriate sampling, one can obtain a linear model of the form,

$$Q(k+1) = Q(k) - S(k+1)\zeta(k) + A(k+1), \quad k \geqslant 0. \tag{9}$$

In this equation $\boldsymbol{Q}$ is the one-dimensional buffer process; $\boldsymbol{S}$ is an i.i.d. sequence of random variables that reflect potential customer service-completions; and let $\boldsymbol{A}$ denotes an i.i.d. sequence representing exogenous arrivals. The stochastic process $\boldsymbol{A}$ is Bernoulli with $\mathsf{P}\{A(k) = 1\} = \alpha(\alpha + \mu)^{-1}$, and $S(k) = 1 - A(k)$, $k \geqslant 1$. The one-dimensional allocation sequence is defined by the *non-idling policy*, $\zeta(k) = \mathbb{I}(Q(k) > 0)$, $k \geqslant 1$.

Suppose now that we wish to approximate a queue with general arrival and service distributions using a similar linear model. Given a sampling increment $T_s > 0$, we consider a model of the form

$$Q(t; x) = x + \sum_{k=1}^{\lfloor t/T_s \rfloor} \bigl[-S(k)\zeta(k) + A(k)\bigr], \quad x \in \mathbb{R},\ t \geqslant 0.$$

We assume that $\boldsymbol{X} = \{X(k) := (S(k), A(k))^{\mathrm{T}}\colon\ k \geqslant 1\}$ is an i.i.d. process, and that the mean and variance of $\boldsymbol{X}$ are of the form

$$\mathsf{E}\bigl[X(k)\bigr] = (\mu, \alpha)^{\mathrm{T}} T_s, \qquad \operatorname{Var}\bigl[X(k)\bigr] = \Sigma T_s,$$

where μ, α again denote service and arrival rates, and $\Sigma \geqslant 0$ is a 2×2 covariance matrix.

For example, suppose that arrivals to the queue are accurately approximated by a delayed renewal process $\boldsymbol{N} = (N(t)\colon\ t \geqslant 0)$. For a given sampling interval $T_s > 0$, define the sampling times $\{t_k = kT_s\colon\ k \geqslant 0\}$. We wish to construct a distribution for the i.i.d. process $\boldsymbol{A}$ so that the random variable $N(t_k)$ has approximately the same distribution as

$$\widetilde{N}(t_k) := \sum_{i=0}^{k} A(i).$$

Consideration of large k suggests the two restrictions,

$$\mathsf{E}\bigl[A(1)\bigr] = \frac{1}{m} T_s \quad \text{and} \quad \operatorname{Var}\bigl[A(1)\bigr] = \frac{\sigma^2}{m^3} T_s,$$

where (m, σ^2) denote the mean and variance of the increment process for $\boldsymbol{N}$. It then follows that the processes $\widetilde{\boldsymbol{N}}$ and $\boldsymbol{N}$ share several asymptotic properties. In particular, the law of large numbers and central limit theorem scalings lead to identical limits.

Approximations of this form may be applied to a general stochastic model with renewal arrival and service statistics, and this leads to the following linear stochastic model. The model (10) was previously considered in [28] where it is called the 2-*parameter model*. The terminology there is motivated by the requirement that the user specify both the mean and variance of the matrices $B(k)$ and vectors $A(k)$, for $k \geqslant 1$.

The controlled random-walk model. The state process $\boldsymbol{Q}$ is piecewise constant on $\mathbb{R}$, with potential jumps at the sampling times $\{t_k\colon\ k \geqslant 1\}$ given by

$$Q(t_{k+1}; x) = Q(t_k; x) + B(k+1)\zeta(k) + A(k+1), \qquad Q(0; x) = x. \tag{10}$$

The following assumptions are imposed on the policy and parameters:

(i) The queue length process satisfies $Q(k) \in \mathsf{X}$, $k \geqslant 0$.

(ii) $\boldsymbol{B}$ is an i.i.d. sequence of $\ell \times \ell$ matrices; and $\boldsymbol{A}$ is an i.i.d. sequence of ℓ-dimensional vectors. It is assumed that the variances are all finite, that $\mathsf{E}[A(k)] = T_s\alpha$, $k \geqslant 1$, and that $\boldsymbol{B}$ is of the specific form,

$$B(k) = -(I - R)M(k), \quad k \geqslant 1,$$

where $\boldsymbol{M}$ is a i.i.d. sequence of diagonal matrices satisfying $\mathsf{E}[M(k)] = T_S M$, and the matrices R, M are given in (5).

(iii) The allocation sequence $\boldsymbol{\zeta}$ is adapted to $(\boldsymbol{Q}, \boldsymbol{A}, \boldsymbol{B})$, and satisfies $\zeta(k) \in \mathsf{U}$, $k \geqslant 0$, where U is defined in (4). A process $\boldsymbol{\zeta}$ satisfying these constraints is called *admissible*.

Note that there is great flexibility in the choice of distribution for $(\boldsymbol{A}, \boldsymbol{B})$. The mean and variances, and perhaps higher-order statistics will be chosen so that the model accurately reflects the behavior of the physical system.

If one wishes to construct an optimal policy using dynamic programming techniques, then it may be convenient to choose distributions that allow $\boldsymbol{Q}$ to be restricted to an integer lattice. Also, it may be physically unnatural to allow non-integer values for $\boldsymbol{\zeta}$. In the numerical results presented in section 6 we impose the following additional restriction,

$$\zeta(k) \in \mathsf{U}_\cdot, \quad k \geqslant 1, \quad \mathsf{U}_\cdot := \mathsf{U} \cap \left\{u \in \mathbb{R}_+^\ell \colon u_i \in \{0, 1\}, 1 \leqslant i \leqslant \ell\right\}. \tag{11}$$

Under the assumptions on the sequences $(\boldsymbol{A}, \boldsymbol{B})$, the fluid model (3) may be justified through scaling the system equations (10) (see, e.g., [9,14]). For $r \geqslant 1$ define

$$q^r(t; x) = \frac{Q(rt; rx)}{r}, \quad x \in \mathsf{X},\ t \geqslant 0. \tag{12}$$

Suppose that the open-loop, constant control is applied, $z(t) = \zeta t$, $t \geqslant 0$, where $\zeta \in \mathsf{U}$ is given. We then have the approximation, for any initial or $x \in \mathsf{X}$, and any time $t \in \mathbb{R}_+$,

$$q^r(t; x) \to x + B\zeta t + \alpha t, \quad r \to \infty \text{ a.s.}$$

2.3. *Workload relaxations*

It is often the case that the *important* dynamics of the network model are captured by a subset of the components of the workload process. The idea is that only some of the constraints on the allocation process are important from the point of view of policy synthesis, and removing less restrictive constraints leads to a far simpler control problem.

Consider first the fluid model:

Workload relaxation for the fluid model. The state process $\hat{\boldsymbol{q}}$ is defined as a *differential inclusion* satisfying,

(i) The state-space constraints

$$\hat{q}(t; x) \in \mathsf{X}, \quad t \geqslant 0.$$

(ii) The rate constraints

$$\frac{\hat{q}(t; x) - \hat{q}(s; x)}{t - s} \in \widehat{\mathsf{V}}, \quad 0 \leqslant s \leqslant t,$$

where the velocity set is given by

$$\widehat{\mathsf{V}} = \left\{ v\colon \langle \xi^i, v \rangle \geqslant -(1 - \rho_i),\ 1 \leqslant i \leqslant n \right\}.$$

(iii) The associated workload process is given by

$$\widehat{w}(t; x) = \widehat{\Xi} \hat{q}(t; x), \quad t \geqslant 0,\ x \in \mathsf{X},$$

where $\widehat{\Xi}$ is the $n \times \ell$ matrix with rows equal to $\{\xi^i\colon 1 \leqslant i \leqslant n\}$.

The dynamics of $\widehat{\boldsymbol{w}}$ are *decoupled* since the workload vectors $\{\xi^i\colon 1 \leqslant i \leqslant n\}$ are linearly independent. This and related properties of the workload process are summarized in the following:

Proposition 2.2. For any feasible trajectory $\hat{\boldsymbol{q}}$, the corresponding workload process satisfies the following:

(i) The workload process $\widehat{\boldsymbol{w}}$ is subject to the decoupled rate constraints,

$$\frac{\mathrm{d}^+}{\mathrm{d}t} \widehat{w}_i(t; x) \geqslant -(1 - \rho_i), \quad 1 \leqslant i \leqslant n. \tag{13}$$

(ii) Define the corresponding idleness process $\boldsymbol{I} = \{I(t) \in \mathbb{R}^n_+\colon t \in \mathbb{R}_+\}$ by

$$I(t) := w(t) + t\delta - w, \quad t \geqslant 0.$$

This is nondecreasing: Setting $\iota(t) = (\mathrm{d}^+/\mathrm{d}t) I(t)$, $t \geqslant 0$, we have

$$\iota_j(t) \geqslant 0, \quad t \in \mathbb{R}_+,\ j = 1, \ldots, n. \tag{14}$$

(iii) The workload process is constrained to the workload space,

$$\widehat{w}(t; x) \in \mathsf{W} := \left\{ \widehat{\Xi} x\colon x \in \mathsf{X} \right\}. \tag{15}$$

The set $\mathsf{W} \subset \mathbb{R}^n_+$ is a positive cone whenever $\mathsf{X} = \mathbb{R}^\ell_+$.

We may also construct a relaxation of the CRW model as follows. The following assumption is imposed to justify a workload model in which variability is consistent across non-idling allocations: Fix $1 \leqslant n \leqslant \ell_r$, and let $\widehat{C}$ denote the $n \times \ell$ matrix obtained by deleting all but the first n rows of C.

$$\begin{aligned}&\text{For any two vectors } \{\zeta^1, \zeta^2\} \subset \mathsf{U}. \text{ satisfying } \widehat{C}\zeta^1 = \widehat{C}\zeta^2,\\ &\text{the distributions of } \big\{\widehat{C}M^{-1}M(k)\zeta^i\colon\ i = 1, 2\big\} \text{ are identical.}\end{aligned} \tag{16}$$

In particular, this implies that the coefficient of variation of $M_{jj}(k)$ is constant over $j \in \mathcal{C}_i$ for any $1 \leqslant i \leqslant n$.

To construct a relaxation of the CRW model we fix a vector $\zeta^\circ \in \mathsf{U}.$ satisfying $\widehat{C}\zeta^\circ = \mathbf{1}$. The dynamics are then defined as follows:

CRW workload model. The workload process $\widehat{\boldsymbol{W}}$ is piecewise constant on $\mathbb{R}$, with potential jumps at times $\{t\colon\ k \geqslant 1\}$, and subject to the recursion

$$\widehat{W}(t_{k+1}) = \widehat{W}(t_k) + E(k+1)\iota(k) - D(k+1), \quad k \geqslant 0. \tag{17}$$

The state process $\widehat{\boldsymbol{W}}$, the idleness process $\boldsymbol{\iota}$, and the processes $\{\boldsymbol{E}, \boldsymbol{D}\}$ are subject to the following constraints:

(i) $\widehat{\boldsymbol{W}}$ is constrained to the workload space W defined in (15).

(ii) The i.i.d. processes $(\boldsymbol{D}, \boldsymbol{E})$ are given by

$$\begin{aligned} D(k) &= \widehat{C}M^{-1}M(k)\zeta^\circ - \widehat{\Xi}A(k),\\ E(k) &= \operatorname{diag}\big(\widehat{C}M^{-1}M(k)\zeta^\circ\big), \quad k \geqslant 1. \end{aligned} \tag{18}$$

(iii) The idleness process $\boldsymbol{\iota}$ is non-negative, and adapted to the processes $(\boldsymbol{D}, \boldsymbol{E})$. A process $\boldsymbol{\iota}$ satisfying these constraints is called *admissible*.

Given the exchangeability assumption (16), we find that this workload model is a true relaxation of the primary model (10). This can provide tremendous advantages in approximating optimal control solutions using dynamic programming techniques.

We have seen that the fluid model may be justified through the scaling (12) of the CRW model, and application of the Law of Large Numbers to the i.i.d. processes $(\boldsymbol{A}, \boldsymbol{B})$. Similarly, one obtains a model with Gaussian disturbances through a Central Limit theorem scaling.

One such scaling is through the sampling interval. Consider a sequence of models, parameterized by the sampling interval $T_s = r^{-1}$. On letting $r \to \infty$, we find that $\boldsymbol{Q}$ converges in distribution to a model with Levy statistics. If the marginal distributions of $(\boldsymbol{A}^r, \boldsymbol{B}^r)$ are chosen appropriately, then the limiting process will be described by a linear fluid model with Gaussian additive disturbance, analogous to the controlled Brownian-motion model described below.

Alternatively, the model can be obtained by considering a sequence of networks parameterized by load. In many specific examples it is shown that this procedure leads to a Gaussian model [7,11,20,24,29,32,33,37,46].

However, in this paper we avoid justification of the existence of a limiting process, and instead *define* a formal stochastic workload model by the introduction of an exogenous Gaussian disturbance.

The definition of the controlled Brownian-motion (CBM) workload model depends upon a non-negative scaling parameter denoted κ. This is introduced to facilitate consideration of sensitivity to variability in the analysis below.

Controlled Brownian-motion workload model. The continuous-time workload process obeys the dynamics:

$$\widehat{W}(t; w, \kappa) = w - (\mathbf{1} - \boldsymbol{\rho})t + I(t) + \sqrt{\kappa} N(t), \qquad \widehat{W}(0) = w \in \mathsf{W}. \tag{19}$$

The state process $\widehat{\boldsymbol{W}}$, the idleness process $\boldsymbol{I}$, and the disturbance process $\boldsymbol{N}$ are subject to the following constraints:

(i) $\widehat{\boldsymbol{W}}$ is constrained to the workload space W defined in (15).

(ii) The stochastic process $\boldsymbol{N}$ is a drift-less n-dimensional Brownian motion, whose instantaneous covariance satisfies $\Sigma > 0$.

(iii) The idleness process $\boldsymbol{I}$ is constrained to be adapted to the Brownian motion $\boldsymbol{N}$, with the simple constraint $I_i(t) - I_i(s) \geqslant 0$ for all $t \geqslant s$, and all $1 \leqslant i \leqslant n$. A process $\boldsymbol{I}$ satisfying these constraints is called *admissible*.

For $\kappa = 1$ we suppress the dependency of $\widehat{W}$ on κ, and write $\widehat{W}(t; w)$ for the process starting from $w \in \mathsf{W}$. We take $\kappa = 1$ unless stated otherwise.

We have painted a fairly broad picture of fluid and stochastic network models. The next section focuses in greater detail on the corresponding control synthesis problem.

3. Optimal control

Throughout the paper we denote by $c : \mathbb{R}^{\ell}_{+} \to \mathbb{R}_{+}$ a *cost function* on buffer-space. We assume throughout that the cost function is piecewise linear.

The cost function $c : \mathbb{R}^{\ell}_{+} \to \mathbb{R}_{+}$ is piecewise linear, convex, vanishes only at the origin, and has the form

$$c(x) = \max\bigl(\langle c^i, x\rangle \colon i = 1, \dots, \ell_c\bigr), \quad x \in \mathbb{R}^{\ell}_{+}, \tag{20}$$

with $c^i \in \mathbb{R}^{\ell}$, $i = 1, \dots, \ell_c$.

We consider in this section various optimal control criteria based on this cost function for the queueing models described in section 2. Many of the results in this section are minor extensions of results from [40,41,43].

We begin with basic definitions and results for the fluid model.

3.1. Optimal control for the fluid model

Allocations for the fluid model are frequently defined through state feedback, $\zeta(t) = f(q(t))$, where $f : \mathsf{W} \to \mathsf{U}$. The *feedback law* f is a special case of a *policy*, which is a fixed set of rules that determine allocations based on observed behavior of the state process. The feedback law f is said to be *stabilizing* if $q(t; x) \to \boldsymbol{\theta}$ as $t \to \infty$, from any initial condition x. A stabilizing feedback law can only exist if $\rho < 1$.

Below are three formulations of *optimal control*.

Time-optimal control. For any initial condition $q(0) = x$, find a control which minimizes

$$T(x) = \min\{t\colon\ q(t; x) = \boldsymbol{\theta}\}.$$

We denote by $T_*(x)$ the infimum over all policies.

Infinite-horizon-optimal control. For any initial condition $q(0) = x$, find a control that minimizes $J(x)$ as defined in (1). We let $J_*(x)$ denote the 'optimal cost', i.e., the infimum over all policies when $q(0) = x$.

Pathwise-optimal control. For any given x, find a control satisfying, for *all* $t \geqslant 0$,

$$c\big(q(t; x)\big) = \min\{c(y)\colon\ y \in \mathsf{X} \cap \{x + t\mathsf{V}\}\}.$$

In addition to optimal policies, we consider the *myopic* (or *greedy)* policy. This is defined as follows: For each $t \geqslant 0$, the allocation rates $\zeta(t)$ are chosen to minimize $(\mathrm{d}^+/\mathrm{d}t)c(q(t; x))$ over all $\zeta \in \mathsf{U}$.

A pathwise-optimal feedback law is simultaneously myopic, time-optimal, and infinite-horizon-optimal.

Suppose that two states $x, y \in \mathsf{X}$ are given with $\langle \xi^i, x\rangle \leqslant \langle \xi^i, y\rangle$ for all $1 \leqslant i \leqslant n$. For any $r > 0$ let $v = r(y - x)$. This velocity vector lies in $\widehat{\mathsf{V}}$, and the feasible trajectory $\hat{q}(t; x) = x + tv$, $0 \leqslant t \leqslant 1/r$, reaches y in $1/r$ seconds. It follows that the minimum time to reach y starting from x is *zero*. This leads to the following terminology.

Effective cost for a workload-relaxation.

(i) The *effective cost* $\bar{c} : \mathsf{W} \to \mathbb{R}_+$ is defined for $w \in \mathsf{W}$ as the value of the linear program

$$\begin{aligned} &\bar{c}(w) = \min r \\ \text{subject to}\quad &r \geqslant \langle c^i, x\rangle, \quad 1 \leqslant i \leqslant \ell_c, \\ &\widehat{\Xi} = w, \\ &x \in \mathsf{X}. \end{aligned} \tag{21}$$

(ii) The region where $\bar{c}$ is *monotone* is denoted W^+. That is,

$$\mathsf{W}^+ := \{w \in \mathsf{W}\colon\ \bar{c}(w') \geqslant \bar{c}(w) \text{ whenever } w' \geqslant w.\}$$

(iii) For any $w \in \mathsf{W}$, the *effective state* $\mathcal{X}^*(w)$ is defined to be any vector $x \in \mathsf{X}$ that minimizes the linear program (21):

$$\mathcal{X}^*(w) = \arg\min_{x \in \mathsf{X}} \big(c(x)\colon\ \widehat{\Xi}x = w\big). \tag{22}$$

(iv) For any $x \in \mathsf{X}$, the optimal exchangeable state $\mathcal{P}^*(x) \in \mathsf{X}$ is defined via

$$\mathcal{P}^*(x) = \mathcal{X}^*\big(\widehat{\Xi}x\big). \tag{23}$$

Since c is piecewise linear, it follows that this is also true for the effective cost

$$\bar{c}(w) = \max_i \langle \bar{c}^i, w \rangle, \tag{24}$$

where $\{\bar{c}^i\} \in \mathbb{R}^n$ are the extreme points obtained in the dual of (21).

The fluid value function, denoted $\hat{J}_*$, is the value function defined for $w \in \mathsf{W}$ by

$$\hat{J}_*(w) = \min \int_0^\infty \bar{c}\big(\hat{w}(s; w)\big)\, \mathrm{d}s,$$

where again the minimum is over all policies. For small values of n, an optimal solution to the workload-relaxation is frequently pathwise-optimal [43].

The following result follows from the fact that the workload relaxation is a true relaxation of the fluid model, in the sense that it is subject to few constraints.

Proposition 3.1. Let J_*, $\hat{J}_*$ denote the value functions for the fluid model and its relaxation, respectively. Then,

$$J_*(x) \geqslant \hat{J}_*\big(\widehat{\Xi}x\big), \quad x \in \mathsf{X}.$$

Bounds in the reverse direction are obtained in [43].

To illustrate the above definitions, again consider the 3-buffer model given in (8). When $\mathsf{X} = \mathbb{R}^3_+$, $c(x) = c^{\mathrm{T}}x$, $x \in \mathsf{X}$, and with $\mu_1 = \mu_3$, the effective cost $\bar{c}(w)$ is the solution to the linear program,

$$\max\, [r_1\mu_1 w_1 + r_2\mu_2 w_2],$$

$$\text{subject to} \quad \begin{bmatrix} 2 & 1 \\ 1 & 1 \\ 1 & 0 \end{bmatrix} r \leqslant \begin{bmatrix} c_1 \\ c_2 \\ c_3 \end{bmatrix}.$$

Figure 3 shows a sublevel set of $\bar{c}$ for the model shown in figure 2 with cost function $c(x) = x_1 + x_2 + x_3$. The effective cost $\bar{c}$ is not monotone since W^+ is a strict subset of W. This is reasonable since reducing workload at machine 2 does not necessarily reduce cost with this cost function. The following proposition summarizes properties of optimal control solutions in this special case. The conclusions all follow from the definitions and [34, theorem 3.4].

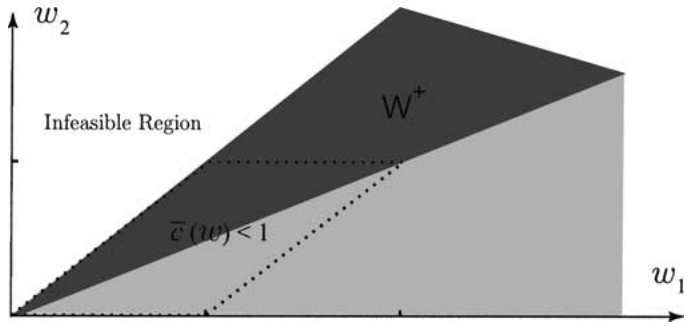

Figure 3. The effective cost for the three buffer model shown in figure 2 is not monotone since feeding the second station conflicts with draining buffer three.

Proposition 3.2. For the fluid model with parameters given in (8), when the cost function $c(x) = x_1 + x_2 + x_3$ is considered for the original network, and with $\mu_1 = \mu_3$, the optimal solution for the second workload-relaxation has the following properties:

(i) $\mathsf{W} = \{w\colon\ 0 \leqslant w_2/w_1 \leqslant 2\rho_2/\rho_1\}$.

(ii) The effective cost is given by $\bar{c}(w) = \max(\mu_2 w_2, \mu_1 w_1 - \mu_2 w_2)$.

(iii) The monotone region for the effective cost is given by

$$\mathsf{W}^+ = \left\{ w\colon\ \frac{\rho_2}{\rho_1} \leqslant \frac{w_2}{w_1} \leqslant 2\frac{\rho_2}{\rho_1} \right\}.$$

(iv) The infinite-horizon-optimal control is described by a linear switching curve,

$$s_*(w_1; 0) = m_* w_1, \quad \text{with } m_* \leqslant \overline{m} := \frac{\rho_2}{\rho_1},$$

so that the optimal paths are constrained to lie in the set $\mathcal{R}_*(0) = \{w\colon\ m_* w_1 \leqslant w_2 \leqslant 2(\rho_2/\rho_1) w_1\}$.

(v) The optimal solution satisfies $(\mathrm{d}/\mathrm{d}t)\widehat{w}_i^*(t) = -(1-\rho_i)$, $i = 1, 2$, whenever $\widehat{w}^*(t)$ lies in the interior $\mathcal{R}_*(0)$.

(vi) If $\rho_2 \leqslant \rho_1$ then $m_* = \overline{m}$ and there exists a pathwise-optimal solution from each initial condition.

(vii) If $\rho_2 \leqslant \rho_1$ and the initial condition $\widehat{w}^*(0) = w$ satisfies

$$\frac{w_2}{w_1} \leqslant \frac{\rho_2}{\rho_1},$$

then a pathwise-optimal solution cannot exist.

Consider, for example, the parameters

$$\mu_1 = \mu_3 = 22, \qquad \mu_2 = 10, \quad \text{and} \quad \alpha^{\mathrm{T}} = (9, 0, 0). \tag{25}$$

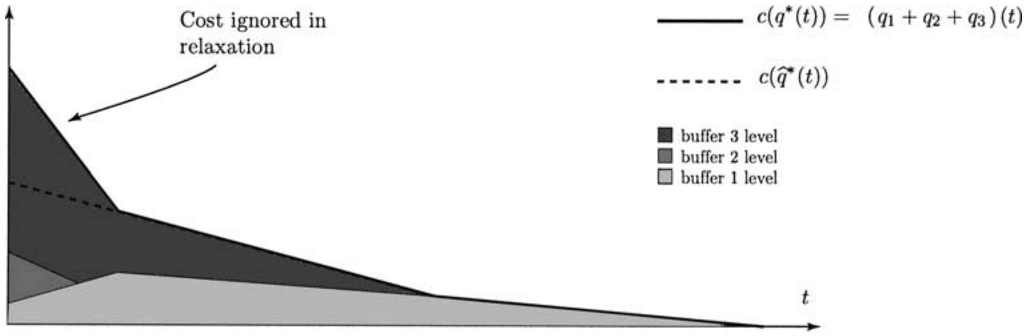

Figure 4. Optimal cost-curves for the second workload-relaxation. The dashed lines show the cost $c(\hat{q}^*(t; x))$ for the optimized one and two-dimensional workload-relaxations. The actual optimal policy incurs a higher cost when compared with the policy for the second workload-relaxation, but this error is bounded in ρ.

We have $\rho_2 = 9/10 > 9/11 = \rho_1$, so the conditions of proposition 3.2(vii) hold, and we conclude that, from certain initial conditions, a pathwise-optimal solution does not exist.

Infinite-horizon-optimal trajectories for the relaxed and original fluid model are illustrated in figure 4. In workload space, the infinite-horizon-optimal policy for $\hat{\boldsymbol{q}}$ is defined by the switching-curve,

$$s_*(w_1; 0) = m_* w_1 = \frac{11}{15} w_1, \quad w_1 > 0.$$

That is, station two is permitted to idle as long as $15w_2 < 11w_1$. In the relaxed model, if the initial workload is below this line, then the optimal state trajectory jumps upward to reach it at time $t = 0+$. Henceforth, the idleness process for the second machine satisfies $(\mathrm{d}/\mathrm{d}t)I_2(t) = 0$ until the upper boundary of the workload space W is reached. The unrelaxed model catches up with the relaxed model after a short transient period, as predicted by [43, theorem 4.1].

We now examine discounted-cost optimal control for stochastic models.

3.2. *Discounted optimal-control*

The discounted optimal-control problem may be formulated for the CRW model, its relaxation, or the CBM workload model.

We begin with the CRW model (10). For a given policy and initial condition $x \in \mathsf{X}$, the associated discounted cost criterion is defined by

$$K(x; \eta) = \int_0^\infty \mathrm{e}^{-\eta s} \mathsf{E}\big[c\big(Q(s; x)\big)\big] \,\mathrm{d}s. \tag{26}$$

The discount parameter satisfies $\eta > 0$, so that K is finite-valued for 'reasonable policies'. The optimal value, denoted $K_*(x; \eta)$, is defined to be the infimum of (26) over all admissible policies, subject to the additional integral constraint (11).

The value function may also be expressed,

$$K_*(x;\eta) = \inf \mathsf{E}\left[\int_0^{\eta^{-1}T} c\big(Q(s;x)\big)\,\mathrm{d}s\right], \quad x \in \mathsf{X}, \tag{27}$$

where T is an exponential random variable with unit mean. The random variable T is assumed to be independent of $(\boldsymbol{A}, \boldsymbol{B})$ in the CRW model.

The value function is defined analogously for either of the stochastic workload models:

$$\widehat{K}_*(w;\kappa,\eta) = \inf \mathsf{E}\left[\int_0^{\eta^{-1}T} \bar{c}\big(\widehat{W}(s;w)\big)\,\mathrm{d}s\right], \quad w \in \mathsf{W}, \tag{28}$$

where T is again a standard exponential random variable, and the infimum is over all admissible idleness processes. The random variable T is assumed to be independent of $(\boldsymbol{D}, \boldsymbol{E})$ in the CRW workload model, and independent of $\boldsymbol{N}$ in the CBM workload model. When $\kappa = 1$ we shall let $\widehat{K}_*(w;\eta)$ denote the value function.

The following result follows from the observation that the optimal control problem for the CRW workload model is a true relaxation of the primary stochastic control problem (27). It is entirely analogous to proposition 3.1.

Proposition 3.3. Let K_*, $\widehat{K}_*$ denote the value functions for the CRW model and its relaxation, respectively. Then,

$$K_*(x;\eta) \geqslant \widehat{K}_*\big(\widehat{\Xi}x;\eta\big), \quad x \in \mathsf{X}.$$

It is likely that the gap $|K_*(x;\eta) - \widehat{K}_*(\widehat{\Xi}x;\eta)|$ may be bounded in heavy traffic, with $\rho \sim 1$ and $\eta \sim 0$, following the bounds obtained in [43, theorems 4.4 and 4.5]. However, this is beyond the scope of the present paper.

The value function is the solution to an associated dynamic programming equation, given here for the workload models:

$$\widehat{K}_*(w;\eta) = \inf \mathsf{E}\left[\int_0^{T_0} \mathrm{e}^{-\eta s}\bar{c}\big(\widehat{W}(s;w)\big)\,\mathrm{d}s + \mathrm{e}^{-\eta T_0}\widehat{K}_*\big(\widehat{W}(T_0;w),\eta\big)\right], \quad w \in \mathsf{W}. \tag{29}$$

The time $T_0 > 0$ is arbitrary in the CBM workload model, and in the CRW workload model we may take $T_0 = t_k$, for any given $k \geqslant 1$.

Moreover, the dynamic programming equation (29) defines the optimal policy, when it exists. Existence is guaranteed for the CRW workload model: given that $\widehat{W}(t_k;w) = w^1$, the optimal idleness rate at time t_k is given by

$$\iota^*(k) = \arg\min\big\{\mathsf{E}\big[\widehat{K}_*\big(w^1 + E(1)u - D(1);\eta\big)\big]\big\}, \tag{30}$$

where the minimum is over all $u \in \mathbb{R}^n_+$ such that the constraint $w^1 + E(1)u - D(1) \in \mathsf{W}$ holds with probability one.

Convexity of the state space and cost function lead to the following properties for $\widehat{K}_*$.

Theorem 3.4. For either the CRW or CBM workload model:

(i) The value function $\widehat{K}_*$ is convex on W. For the CBM workload model, $\widehat{K}_*$ is also monotone.

(ii) The value function is linearly bounded:

$$\limsup_{\|w\|\to\infty}\left(\frac{\widehat{K}_*(w;\eta)}{\|w\|}\right) < \infty.$$

(iii) Its gradient is bounded, and we have the explicit bound

$$\widehat{K}_*\big(w + w^1;\eta\big) - \widehat{K}_*(w;\eta) \leqslant \eta^{-1}\bar{c}\big(w^{-1}\big), \quad w, w^1 \in \mathsf{W}.$$

(iv) The value function scales as follows:

$$\widehat{K}_*(w;\kappa,\eta) = \kappa\widehat{K}_*\big(\kappa^{-1}w;1,\eta\kappa\big), \quad w \in \mathsf{W},\ \kappa,\eta > 0.$$

Proof. We first establish convexity. Let $w^0, w^1 \in \mathsf{W}$, $\vartheta \in [0,1]$, and set $w^2 = \vartheta w^0 + (1-\vartheta)w^1$. Let $I(t;w^0)$, $I(t;w^1)$ be admissible idleness processes corresponding to the initial conditions w^0 and w^1, respectively. The following idleness process is then admissible from the initial condition w^2:

$$I\big(t;w^2\big) := \vartheta I^*\big(t;w^0\big) + (1-\vartheta)I^*\big(t;w^1\big), \quad t \geqslant 0,$$

and we then have $\widehat{W}(t;w^2) = \vartheta\widehat{W}^*(t;w^0) + (1-\vartheta)\widehat{W}^*(t;w^1)$ for all $t \geqslant 0$. From convexity of $\bar{c}$ we must then have $\bar{c}(\widehat{W}(t;w^2)) \leqslant \vartheta\bar{c}(\widehat{W}^*(t;w^0)) + (1-\vartheta)\bar{c}(\widehat{W}^*(t;w^1))$, and on integrating over t one obtains the inequality

$$\widehat{K}_*\big(w^2;\eta\big) \leqslant \widehat{K}_2\big(w^2;\eta\big) \leqslant \vartheta\widehat{K}_0\big(w^0;\eta\big) + (1-\vartheta)\widehat{K}_1\big(w^1;\eta\big),$$

where $\widehat{K}_i(w^i;\eta)$ is the discounted cost incurred using the idleness process $I(t;w^i)$. Taking the infimum over all idleness processes establishes convexity, and this completes the proof of (i).

We now show that $\widehat{K}_*$ is monotone for the CBM workload model. Suppose that $w, w' \in \mathsf{W}$, with $w' \geqslant w$. Then any trajectory $\{\widehat{W}(t;w')\colon 0 < t < \infty\}$ starting at w' is also feasible for the process starting at w. Monotonicity immediately follows.

The proof of (ii) is obvious from the homogeneity of the cost function and network dynamics.

The proof of (iii) is similar to the proof of convexity. Let $w \in \mathsf{W}$ be arbitrary, and let $I(t;w)$ denote any admissible idleness process. Then, this idleness process is also admissible for the process starting from the initial condition $w + w^1$, and we would then have

$$\begin{aligned}\widehat{W}\big(t;w+w^1\big) &= \widehat{W}(t;w) + w^1,\\ \bar{c}\big(\widehat{W}\big(t;w+w^1\big)\big) &\leqslant \bar{c}\big(\widehat{W}(t;w)\big) + \bar{c}\big(w^1\big), \quad t \geqslant 0,\end{aligned}$$

where the inequality follows from the fact that $\bar{c}$ is a norm, and hence the triangle inequality holds. Taking expectations, integrating, and infimizing over $\boldsymbol{I}$ gives

$$\widehat{K}_*\big(w + w^1; \eta\big) \leqslant \widehat{K}_*(w; \eta) + \int_0^\infty \mathrm{e}^{-\eta s}\big[\bar{c}\big(w^1\big)\big]\,\mathrm{d}s = \widehat{K}_*(w; \eta) + \eta^{-1}\bar{c}\big(w^1\big),$$

which is the desired bound.

The identity (iv) is a direct consequence of lemma 5.1 below. □

3.3. *Average-cost-optimal control*

We now turn to the average-cost optimization problem described in the introduction. We focus on workload models throughout this subsection. The cost criterion is given by

$$\widehat{K}_*(w) := \inf\left(\limsup_{T\to\infty} \mathsf{E}\left[\frac{1}{T}\int_0^T \bar{c}\big(\widehat{W}(t; w)\big)\,\mathrm{d}t\right]\right),$$

where the infimum is over all policies, and $w \in \mathsf{W}$ is the initial condition. Under our standing assumption that $\Sigma > 0$, it may be shown that $\widehat{K}_*(w)$ is independent of $w \in \mathsf{W}$ (see, e.g., [42]).

An optimal policy is defined by a *relative value function* $h_*: \mathsf{W} \to \mathbb{R}$. This may be constructed as follows:

$$\gamma_*(\kappa) := \limsup_{\eta\downarrow 0} \eta\widehat{K}_*(\boldsymbol{\theta}; \kappa, \eta), \tag{31}$$

$$h_*(w; \kappa) := \liminf_{\eta\downarrow 0}\big(\widehat{K}_*(w; \kappa, \eta) - \widehat{K}_*(\boldsymbol{\theta}; \kappa, \eta)\big), \quad w \in \mathsf{W}. \tag{32}$$

When $\kappa = 1$, and there is no risk of ambiguity, we omit the dependency on this parameter, writing $h_*(w) = h_*(w; 1)$, $w \in \mathsf{W}$.

The pair (h_*, γ_*) solves a dynamic programming equation. The proof of proposition 3.5 follows as in [42, theorem 1.7] (see also [39, theorem 7.2]).

Proposition 3.5. Suppose that $\rho < 1$ for the CRW or CBM workload model. Then, the constant γ_* is finite, and equal to the optimal average cost $\widehat{K}_*(w)$ for any $w \in \mathsf{W}$. Moreover, we have for any policy and any $T_0 > 0$,

$$h_*(w) = \inf \mathsf{E}\left[\int_0^{T_0} \big(\bar{c}\big(\widehat{W}(s; w)\big) - \gamma_*\big)\,\mathrm{d}s + h_*\big(\widehat{W}(T_0; w)\big)\right], \quad T_0 \geqslant 0, \tag{33}$$

where the infimum is with respect to all admissible controls $\boldsymbol{I}$. In the CRW workload model, T_0 is restricted to the sampling times $\{t_k\}$.

For the CRW workload model an optimal policy is obtained exactly as in the discounted case: given that $\widehat{W}(t_k; w) = w^1$, the optimal idleness rate is given by

$$\iota^*(k) = \arg\min\big\{\mathsf{E}\big[h_*\big(w^1 + E(1)u - D(1)\big)\big]\big\}, \tag{34}$$

where the minimum over $u \in \mathbb{R}^n_+$ is constrained as in the discounted case (30).

Moreover, under this policy the relative value function solves the *Poisson equation*

$$\mathsf{E}\big[h_*\big(\widehat{W}^*(t_k; w) + E(k+1)\iota^*(k) - D(k+1)\big) \mid \widehat{W}^*(t_k; w) = w^1\big] = -\bar{c}\big(w^1\big) + \gamma^*,$$
$$w^1 \in \mathsf{W}.$$

For details see [39].

Theorem 3.6 provides bounds on h_* analogous to the previous bounds obtained on $\widehat{K}_*$.

Theorem 3.6. Suppose that $\rho < 1$ for the CRW or CBM workload model. Then,

(i) The relative value function h_* is convex on W. For the CBM workload model, h_* is also monotone.

(ii) The relative value function is quadratically bounded:

$$\limsup_{\|w\|\to\infty}\left(\frac{h_*(w)}{\|w\|^2}\right) < \infty.$$

(iii) Its gradient is linearly bounded:

$$\sup\left(\frac{\|\nabla h_*(w)\|}{\|w\|+1}\right) < \infty,$$

where the supremum is over all $w \in \mathsf{W}$, and all subgradients at w.

(iv) The relative value function scales to the fluid value function: as $n \to \infty$,

$$\frac{1}{n^2}h_*(nw) \to \hat{J}_*(w), \quad w \in \mathsf{W}.$$

Proof. Result (i) follows from proposition 3.4(i). The remaining results are proved for countable state-space queueing models in [39], and the same ideas may be used in these continuous state-space models. □

4. Affine approximations

Affine policies were introduced in [39] as a method of translating a policy for the fluid model to its stochastic counterpart. In this section we provide general conditions under which an optimal switching-curve exists, and can be well approximated by an affine policy. The asymptotic affine shift has a simple, explicit form.

To simplify the development we henceforth concentrate on the CBM workload model in two dimensions. This restriction and further assumptions are listed here:

(A1) We restrict to a two-dimensional workload-relaxation, i.e., $\mathsf{W} \subset \mathbb{R}^2$. We further assume that the cost function $\bar{c}$ given in (19) is monotone in w_1, and that the monotone region W^+ has non-empty interior.

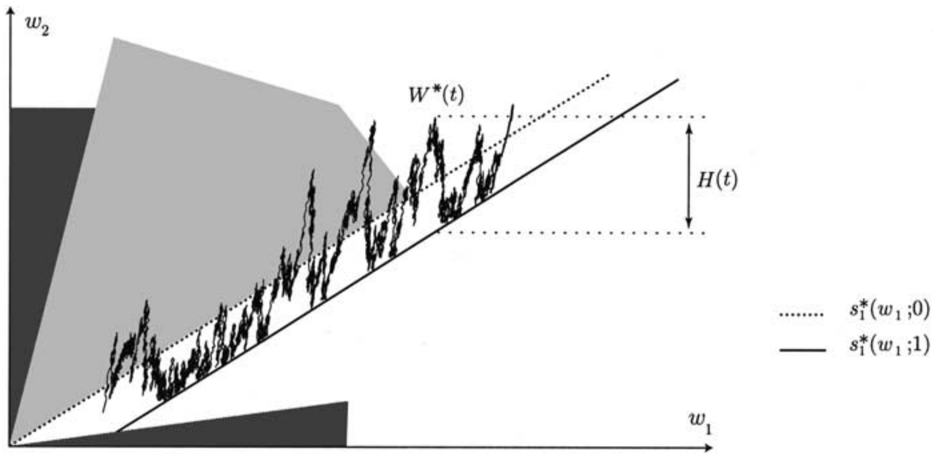

Figure 5. The height process is positive recurrent in case II*.

(A2) $\rho < 1$, so that a solution (h_*, γ_*) exists, solving the average-cost optimal control equation (33) for the model (19) with $\kappa = 1$. It is assumed that $h_*: \mathsf{W} \to \mathbb{R}_+$ is C^1 on W.

(A3) For each $w \in \mathsf{W}$, there exists an admissible idleness process $\{I^*(t; w): t \in \mathbb{R}\}$ that achieves the infimum in (33). The optimal process W is a strong Markov process.

The range of possible behavior in this two-dimensional setting is limited, but in this special case the impact of variability on optimal control solutions is most transparent.

The next result establishes basic structure of average-cost optimal policies under these assumptions. Figure 5 illustrates a switching-curve for a two dimensional model as described in proposition 4.1.

Proposition 4.1. Suppose that (A1)–(A3) hold, and define

$$\mathcal{R}_* := closure\left(\mathsf{W}^+ \cup \left\{w \in \mathsf{W}\colon \frac{\partial}{\partial w_2} h_*(w) > 0\right\}\right). \tag{35}$$

The average-cost optimal solution has the following properties:

(i) The optimal process satisfies $\widehat{W}^*(t; w) \in \mathcal{R}_*$ for all $w \in \mathsf{W}$, $t > 0$, and

$$\int_0^\infty \mathbb{I}\{\widehat{W}^*(t; w) \in \text{interior}(\mathcal{R}_*)\}\, \mathrm{d}I^*(t; w) = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad \text{a.s.}$$

(ii) The optimal region $\mathcal{R}_*$ is defined by an (optimal) switching curve $s_*: \mathbb{R}_+ \to \mathbb{R}$ with

$$\mathcal{R}_* = \{w \in \mathsf{W}\colon w_2 \geqslant s_*(w_1)\}.$$

(iii) If $\bar{c}$ is monotone then $\mathcal{R}_* = \mathsf{W}^+ = \mathsf{W}$.

Proof. The fact that $\mathcal{R}_*$ is a constraint set for $\widehat{\boldsymbol{W}}^*$ follows from the dynamic programming equations (33).

The existence of s_* is a consequence of the convexity of the relative value function h_*; the cost function $\bar{c}$; and the state space W. In particular, convexity of h_* implies that for any $w \in \mathsf{W}$, if $(\partial/\partial w_2)h(w) > 0$, then $(\partial/\partial w_2)h(w + b\mathrm{e}^2) > 0$ for all $b \geqslant 0$, provided $w + b\mathrm{e}^2 \in \mathsf{W}$. This and a similar property for $\bar{c}$ implies that the lower boundary of $\mathcal{R}_*$ may be constructed as a function s_* as claimed.

Part (iii) is the definition of monotonicity. □

Assumption (A1) is imposed so that we may restrict attention to a single switching-curve. If the monotonicity assumption is violated then there may be a second switching-curve that determines the optimal idleness process $\boldsymbol{I}_1^*$. All of our conclusions carry over to this more general setting with only notational changes.

In developing structural properties of optimal control solutions there are three cases to consider, according to the structure of the effective cost:

- *Case I. The effective cost $\bar{c}$ is monotone on* W, i.e., $\mathsf{W}^+ = \mathsf{W}$. In [43, theorems 3.4] it is argued that in this case the fluid-model relaxation admits a pathwise-optimal solution from any initial condition: the optimal trajectories evolve in the set $\mathsf{W}^+ = \mathsf{W}$ for all $t > 0$. When $\kappa > 0$ then we can once again conclude that the optimal solution is trivial, in the sense that $\widehat{\boldsymbol{W}}$ is pointwise minimal (see [43, theorems 4.3–4.5]).
- *Case II. The effective cost $\bar{c}$ is not monotone on* W*, and* $\mathbf{1} - \boldsymbol{\rho} \in \mathsf{W}^+$. In this case, when $\kappa = 0$ the two-dimensional workload model again admits a pathwise-optimal solution from any initial condition, and the optimal trajectories evolve in W^+ (again, see [43, theorem 3.4]). However when $\kappa > 0$, the model (19) does not admit a pathwise-optimal solution from *any* initial condition.

Throughout much of the paper we restrict attention to the strengthened version,

- *Case II*. The vector* $(1 - \rho_1, 1 - \rho_2)^{\mathrm{T}}$ *lies in the interior of* W^+.
- *Case III. The effective cost $\bar{c}$ is not monotone on* W*, and* $\mathbf{1} - \boldsymbol{\rho} \notin \mathsf{W}^+$. In this case, pathwise optimality cannot hold for arbitrary initial conditions, even when $\kappa = 0$. Optimal trajectories for the fluid model evolve in a closed positive cone $\mathcal{R}_*(0)$ that is strictly larger than W^+. It is of the form

$$\mathcal{R}_*(0) = \{w \in \mathbb{R}_+^2\colon\ w_2 \geqslant s_*(w_1; 0)\},$$

where the function $s_* : \mathbb{R} \to \mathbb{R}$ is linear. Again, when $\kappa > 0$, the model (19) does not admit a pathwise-optimal solution from *any* initial condition.

In cases II or III the lower boundary of W^+ is different than the lower boundary of W. We let m_* denote its slope, and we let $\bar{c}^+, \bar{c}^- \in \{\bar{c}^i\}$, denote the cost vectors that define this boundary. Consequently,

$$\begin{aligned} &\bar{c}_1^+ \geqslant 0, \quad \bar{c}_2^+ \geqslant 0, \quad \bar{c}_1^- \geqslant 0, \\ &\bar{c}_2^- < 0, \quad \text{and} \\ &\bar{c}\big((1, m_*)^{\mathrm{T}}\big) = \bar{c}_1^+ + m_*\bar{c}_2^+ = \bar{c}_1^- + m_*\bar{c}_2^-. \end{aligned} \tag{36}$$

Under (Al) we make the following definition:

Affine policies. A policy for (19) is called *affine* if the controlled model is a reflected Brownian motion in an affine domain of the form

$$\mathcal{R}(\kappa) = \{w \in \mathsf{W}\colon\ w_2 \geqslant m w_1 + \kappa d\},$$

with $d, m \in \mathbb{R}$. We let $\widehat{W}(t; w, \kappa, d)$ denote the resulting workload process initialized at $\widehat{W}(0) = w \in \mathsf{W}$.

If the initial condition w lies outside of $\mathcal{R}(\kappa)$, then the process moves instantaneously to this set, so that

$$\widehat{W}(0^{+}; w, \kappa, d) = (w_1, m w_1 + \kappa d)^{\mathrm{T}}.$$

In much of what follows we restrict to the affine policy given by

$$s(w; \kappa, d) = m_* w - \kappa d, \quad w \in \mathbb{R}, \tag{37}$$

where $m_* < \infty$ defines the lower boundary of W^{+}. This is an *affine translation* of the myopic policy for the fluid model. In case II, the policy (37) is an affine translation of the optimal policy $s_*(\cdot; 0)$.

4.1. Some auxiliary processes

In our approximation of the optimal policy using an affine policy we consider a further relaxation of the CBM model, called the *unconstrained process*. This and an associated *height process* are illustrated in figure 5.

The unconstrained process. For a given initial condition $w \in \mathsf{W}$,

(i) The *unconstrained* process $\boldsymbol{X}$ is the CBM model with state space $\mathbb{R}^2$, and cost function $\bar{c}_X(w) := \max(\bar{c}^{+} \cdot w, \bar{c}^{-} \cdot w)$, $w \in \mathbb{R}^2$. The associated value function is given by

$$\widehat{K}^{X}_{*}(w; \kappa, \eta) = \inf \mathsf{E}\left[\int_0^{\eta^{-1}T} \bar{c}_X\big(X(t; w, \kappa)\big)\, \mathrm{d}t\right], \quad w \in \mathbb{R}^2, \tag{38}$$

where the infimum is over all admissible policies.

(ii) For a given switching-curve $s\colon \mathbb{R} \to \mathbb{R}$ defining the idleness process $\boldsymbol{I}_2$ for $\boldsymbol{X}$, the *height process* is defined by $H(t) := X_2(t) - s(X_1(t))$, $t \geqslant 0$.

We restrict primarily to the affine policy (37). In this special case, the height process is the one-dimensional reflected Brownian motion

$$\mathrm{d}H(t) = -\delta_H + \mathrm{d}I(t) + N_H(t), \tag{39}$$

whose drift and instantaneous covariance are given by

$$\delta_H = (\boldsymbol{1} - \boldsymbol{\rho}) \cdot (-m_*, 1)^{\mathrm{T}}, \qquad \sigma^2_H(\kappa) = \kappa \sigma^2_H(1) = \kappa(-m_*, 1)\Sigma(-m_*, 1)^{\mathrm{T}}. \tag{40}$$

In case II, the height process (39) is recurrent since $\delta_H \geqslant 0$. We have positive recurrence in case II* since this inequality is then strict (see lemma A.2).

The unconstrained process may be constructed by scaling the CBM model: for each $r \geqslant 0$ define,

$$\begin{aligned} \mathsf{W}_r &= \{\mathsf{W} - r(1, m_*)^{\mathrm{T}}\} \subset \mathbb{R}^2; \\ \bar{c}_r(w) &= c\big(w + r(1, m_*)^{\mathrm{T}}\big) - r\bar{c}\big((1, m_*)^{\mathrm{T}}\big), \quad w \in \mathsf{W}_r. \end{aligned} \tag{41}$$

Suppose that $\bar{c}$ is not monotone, which implies that $(1, m_*)^{\mathrm{T}}$ lies in the interior of W. It then follows that

$$\mathsf{W}_r \uparrow \mathbb{R}^2 \quad \text{and} \quad \bar{c}_r(w) \downarrow \bar{c}_X(w), \quad r \to \infty,\ w \in \mathbb{R}^2.$$

For each $r \geqslant 0$, one can pose an optimization problem using this cost function and state space. We see below that, as $r \to \infty$, the values converge rapidly to the value of the analogous optimization problem for the process $\boldsymbol{X}$ with state space $\mathbb{R}^2$, and cost function $\bar{c}_X$.

These results are most easily developed in the discounted case that we investigate next.

4.2. Discounted optimal control

We show in this section that the optimal switching-curve s_* is asymptotically affine. No approximation is required for the unconstrained process.

Proposition 4.2. Suppose that (A1)–(A3) are satisfied. The following then hold when $\kappa = 1$ and $\eta > 0$ is arbitrary, regardless of network load, in cases II or III:

(i) There is an affine shift parameter $d_*(\eta) \geqslant 0$ such that the optimal policy for $\boldsymbol{X}$ is defined by the switching curve (37).

(ii) The value function $\widehat{K}_*^X$ is C^1 on $\mathbb{R}^2$, and may be expressed as follows: for $w \in \mathbb{R}^2$ satisfying $w_2 < m_* w_1 - d_*$ we have

$$\widehat{K}_*^X(w; \eta) = \widehat{K}_*^X\big((w_1, m_* w_1 - d_*)^{\mathrm{T}}; \eta\big).$$

For $w \in \mathbb{R}^2$ satisfying $w_2 \geqslant m_* w_1 - d_*$ and $w_2 \neq m_* w_1$,

$$\widehat{K}_*^X(w; \eta) = \eta^{-1}\bar{c}_X(w) - \eta^{-2}(\mathbf{1} - \boldsymbol{\rho})^{\mathrm{T}}\nabla\bar{c}_X(w) + g(w_2 - m_* w_1 + d_*)$$

$$\text{where} \quad g(r) = \begin{cases} A_+ \mathrm{e}^{\varrho + r} + A_- \mathrm{e}^{\varrho - r}, & r < d_*, \\ B_- \mathrm{e}^{\varrho - r}, & r > d_*, \end{cases} \quad r \in \mathbb{R}_+.$$

The coefficients $\{A_+, A_-, B_-\}$ are real constants, and $\varrho_- < 0 < \varrho_+$ are the roots of the quadratic equation

$$\frac{1}{2}\sigma_H^2 \varrho^2 - \delta_H \varrho - \eta = 0.$$

(iii) The optimal idleness process $\boldsymbol{I}_X^*$ is unique.

Proof. To see (i), fix $w^0 \in \mathbb{R}^2$, and define $w^r = w^0 + r(1, m_*)^{\mathrm{T}}$, $r \in \mathbb{R}$. If $X(t; w^0)$ is any feasible trajectory starting from w^0, then

$$X\big(t; w^r\big) := X\big(t; w^0\big) + r(1, m_*)^{\mathrm{T}}, \quad t \geqslant 0,$$

is a feasible trajectory starting from w^r. Moreover, by definition of the cost function,

$$\bar{c}_X\big(X\big(t; w^r\big)\big) = \bar{c}_X\big(X\big(t; w^0\big)\big) + r\bar{c}\big((1, m_*)^{\mathrm{T}}\big), \quad t \geqslant 0,\ r \in \mathbb{R}.$$

Optimizing over all feasible $\boldsymbol{X}$, it follows that

$$\widehat{K}_*^X\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) = \widehat{K}_*^X\big(w^0; \eta\big) + r\eta^{-1}\bar{c}\big((1, m_*)^{\mathrm{T}}\big), \quad t \geqslant 0,\ r \in \mathbb{R},\ \eta > 0. \tag{42}$$

An optimal policy for $\boldsymbol{X}$ is defined by the switching-curve,

$$s_*^X(w_1; \eta) = \sup\left\{w_2 \leqslant m_* w_1\colon \frac{\partial}{\partial w_2}\widehat{K}_*^X\big(w_1, w_2)^{\mathrm{T}}; \eta\big) = 0\right\}, \quad w_1 \in \mathbb{R}.$$

This function is affine since $(\partial/\partial w_2)\widehat{K}_*^X(w^0 + r(1, m_*)^{\mathrm{T}}; \eta)$ does not depend upon r. This proves (i).

The expression for the value function in (ii) also follows from the scaling property (42). The constants $\{\varrho_+, \varrho_-\}$ are chosen so that g solves the following eigenfunction equation for the height process:

$$\mathcal{A}_H g = -\delta_H g' + \frac{1}{2}\sigma_H^2(\kappa)g'' = \eta g.$$

The coefficients $\{A_+, A_-, B_-\}$ are determined by the constraint that $\widehat{K}_*$ is C^1.

The value function $\widehat{K}_*^X$ satisfies $(\mathrm{d}^2/\mathrm{d}w_2^2)\widehat{K}_*^X(w; \eta) > 0$ for all $w \in \mathbb{R}^2$ satisfying $w_2 > m_* w_1 - d_*$. This observation together with the dynamic programming equations may be used to establish (iii). □

The next result establishes solidarity between $\widehat{\boldsymbol{W}}$ and $\boldsymbol{X}$. Let $\widehat{\mathsf{C}} \subset \mathsf{W}$ denote the region on which $\bar{c}$ and $\bar{c}_X$ agree:

$$\widehat{\mathsf{C}} := \big\{w \in \mathsf{W}\colon \bar{c}(w) = \bar{c}_X(w)\big\}.$$

Proposition 4.3. Under assumptions (A1)–(A3) we have

(i) Define $\tau := \inf(t \geqslant 0\colon X^*(t; w^0) \in \widehat{\mathsf{C}}^c)$, where $\boldsymbol{X}^*$ is the optimal unconstrained process. Then, there exists $D_0 > 0$ such that,

$$\mathsf{P}\big\{\tau < \eta^{-1}T \mid X^*(0) = w^0 + r(1, m_*)^{\mathrm{T}}\big\} = \mathrm{O}\big(\mathrm{e}^{-D_0 r}\big), \quad r \geqslant 0,$$

where the bound is uniform over $w^0 \in \widehat{\mathsf{C}}$ satisfying $w_2^0 \leqslant m_* w_1^0$.

(ii) For all $w^0 \in \mathsf{W}$,

$$\widehat{K}_*^X\big(w^0; \eta\big) \leqslant \widehat{K}_*\big(w^0; \eta\big).$$

(iii) Suppose that the affine policy s_*^X is applied to the CBM model $\widehat{\boldsymbol{W}}$, and let $\widehat{K}(w;\eta)$ denote the resulting value function. Then, for all $w^0 \in \mathsf{W}$, $r \geqslant 0$,

$$\begin{aligned}\widehat{K}_*^X\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) &\leqslant \widehat{K}\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big)\\ &\leqslant \widehat{K}_*^X\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) + \mathrm{O}\big(\big(\|w^0\| + 1\big)\mathrm{e}^{-D_0 r}\big),\end{aligned}$$

where D_0 is as in (i).

(iv) For all $w^0 \in \mathsf{W}$,

$$\big\{\widehat{K}_*\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) - r\eta^{-1}\bar{c}\big((1, m_*)^{\mathrm{T}}\big)\big\} \downarrow \widehat{K}_*^X\big(w^0; \eta\big), \quad r \uparrow \infty.$$

Proof. The proof of (i) may be translated to a bound on the height process, which is a one-dimensional RBM on $\mathbb{R}_+$ when $\boldsymbol{X}$ is defined by an affine policy (see figure 5). We omit the details.

The inequality (ii) is a consequence of two observations: (1) $\bar{c}_X(w) \geqslant \bar{c}(w)$ for all $w \in \mathsf{W}$; and (2) the process $\boldsymbol{X}$ is subject to fewer constraints than $\widehat{\boldsymbol{W}}$.

To show (iii) we apply the following identity for the workload process under the affine policy:

$$\begin{aligned}\widehat{K}\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) &= \mathsf{E}\bigg[\int_0^{T\eta^{-1}} \bar{c}\big(\widehat{W}\big(t; w^0 + r(1, m_*)^{\mathrm{T}}\big)\big)\,\mathrm{d}t\bigg]\\ &= \mathsf{E}\bigg[\int_0^{\tau\wedge T\eta^{-1}} \bar{c}_X\big(X^*\big(t; w^0 + r(1, m_*)^{\mathrm{T}}\big)\big)\,\mathrm{d}t\bigg]\\ &\quad + \mathsf{E}\bigg[\mathbb{I}\{\tau \geqslant T\eta^{-1}\}\int_\tau^{T\eta^{-1}} \bar{c}\big(\widehat{W}\big(t; w^0 + r(1, m_*)^{\mathrm{T}}\big)\big)\,\mathrm{d}t\bigg].\end{aligned}$$

An application of the Cauchy–Schwartz inequality to the right-hand side, combined with the bound in (i), implies the desired upper bound.

To see (iv), observe that by theorem 3.4(iii) and proposition 4.2,

$$\begin{aligned}&\frac{\mathrm{d}}{\mathrm{d}r}\widehat{K}_*\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) \leqslant \eta^{-1}\bar{c}\big((1, m_*)^{\mathrm{T}}\big),\\ &\frac{\mathrm{d}}{\mathrm{d}r}\widehat{K}_*^X\big(w^0 + r(1, m_*)^{\mathrm{T}}; \eta\big) = \eta^{-1}\bar{c}_X\big((1, m_*)^{\mathrm{T}}\big) = \eta^{-1}\bar{c}\big((1, m_*)^{\mathrm{T}}\big).\end{aligned}$$

This implies that the left-hand side in (iv) is decreasing in r, and convergence follows from (ii) and (iii). □

Given these results, it is not surprising that the optimal policy for $\widehat{\boldsymbol{W}}$ may be approximated by an affine policy.

Theorem 4.4. Under (A1)–(A3) we have

(i) $\limsup_{w_1\to\infty}(s_*(w_1) - s_*^X(w_1)) = 0$.

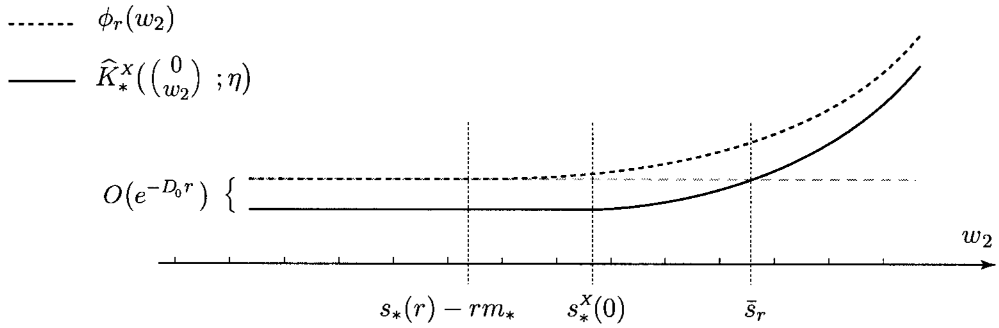

Figure 6. Normalized value function plotted along the w_2 direction. The bound $\bar{s}_r$ decreases to $s^X_*(0) = -d_*$.

(ii) For any $w^0 \in \mathsf{W}$,

$$I^*\big(\cdot;\, w^0 + r(1, m_*)^{\mathrm{T}}\big) \xrightarrow{\mathrm{w}} I^*_X\big(\cdot;\, w^0\big), \qquad r \to \infty,$$

where $\boldsymbol{I}^*_X$ denotes the idleness process for $\boldsymbol{X}^*$. The convergence is in distribution, in the space $D([0, \infty), \mathbb{R}^2_+)$ (see [5]).

Proof. Define for $r \geqslant 0$,

$$\phi_r(w_2) = \widehat{K}_*\big((r, w_2 + rm_*)^{\mathrm{T}};\, \eta\big) - r\eta^{-1}\bar{c}\big((1, m_*)^{\mathrm{T}}\big).$$

This is defined for $w_2 \in \mathbb{R}$ satisfying $(r, w_2{+}rm_*)^{\mathrm{T}} \in \mathsf{W}$. Hence the domain of definition is a closed interval containing the origin, and ϕ_r is convex and increasing as a function of w_2 on this domain. From proposition 4.3 we have, for any $w_2 \in \mathbb{R}$,

$$\phi_r(w_2) \downarrow \widehat{K}^X_*\big((0, w_2)^{\mathrm{T}}\big), \qquad r \uparrow \infty.$$

The following will provide an upper bound on s_*:

$$\bar{s}_r := \min\big\{w_2\colon\ \widehat{K}^X_*\big((0, w_2)^{\mathrm{T}}\big) \geqslant \phi_r\big(s_*(r) - rm_*\big)\big\}.$$

A typical plot is shown in figure 6, from which it is easily seen that $s_*(r) \leqslant rm_* + \bar{s}_r$ for all $r \geqslant 0$, and $\bar{s}_r \downarrow s^X_*(0) = -d_*$. Consequently, we have $\limsup_{w_1\to\infty}(s_*(w_1) - s^X_*(w_1)) \leqslant 0$. We now use this bound to prove (ii), which will also complete the proof of (i).

It is simplest to translate the state $w^0 + r(1, m_*)^{\mathrm{T}}$ to w^0 using (41). For each $r \geqslant 0$ we consider the optimal process $\widehat{\boldsymbol{W}}^*_r$ with initial condition w^0, with $\boldsymbol{N}$ and $\boldsymbol{X}^*$ fixed throughout. It follows from the asymptotic upper bound obtained for s_* that

$$\limsup_{r\to\infty} I^*_r\big(t;\, w^0\big) < I^*_X\big(t;\, w^0\big), \qquad t \geqslant 0.$$

Since $\boldsymbol{I}^*_r$ is also an increasing function of t for each $r \geqslant 0$, we may conclude that the associated distributions are tight.

Let $\boldsymbol{I}^\circ$ be a representation of a weak limit: for some sequence $\{r_i\}$ we have,

$$\left\{N(\cdot), I^*_{r_i}\left(\cdot; w^0\right)\right\} \xrightarrow{\text{w}} \left\{N(\cdot), I^\circ\left(\cdot; w^0\right)\right\}, \quad i \to \infty.$$

The processes $(\boldsymbol{N}, \boldsymbol{I}^\circ\}$ are defined on a common probability space, and $\boldsymbol{I}^\circ$ is adapted to $\boldsymbol{N}$. We can thus define,

$$X^\circ\left(t; w^0\right) = w^0 - (\mathbf{1} - \boldsymbol{\rho})t + I^\circ\left(t; w^0\right) + N(t), \quad t \geqslant 0.$$

We can establish uniform integrability to prove the following limit,

$$\lim_{i\to\infty} \mathsf{E}\left[\bar{c}_{r_i}\left(\widehat{W}^*_{r_i}\left(\eta^{-1}T; w^0\right)\right)\right] = \mathsf{E}\left[\bar{c}_X\left(X^\circ\left(\eta^{-1}T; w^0\right)\right)\right].$$

From elementary properties of exponential random variables, we also have, for any $r \geqslant 0$,

$$\mathsf{E}\left[\bar{c}_r\left(\widehat{W}^*_r\left(\eta^{-1}T; w^0\right)\right)\right] = \eta \widehat{K}^r_*\left(w^0; \eta\right),$$

and hence from proposition 4.3,

$$\lim_{i\to\infty} \mathsf{E}\left[\bar{c}_{r_i}\left(\widehat{W}^*_{r_i}\left(\eta^{-1}T; w^0\right)\right)\right] = \eta \widehat{K}^X_*\left(w^0; \eta\right).$$

We conclude that $\boldsymbol{I}^\circ$ is an optimal allocation for $\boldsymbol{X}$, satisfying $I^\circ(t; w^0) < I^*(t; w^0)$ for all $t \in \mathbb{R}_+$. This is only possible if the processes $\boldsymbol{I}^\circ$ and $\boldsymbol{I}^*$ are identical (see proposition 4.2). □

4.3. Average cost optimal control

The construction of the relative value function h_* given in (32) is based upon the value functions $\{\widehat{K}_*(\cdot; \eta)\colon \eta > 0\}$. As η tends to zero one might also consider the sequence of optimal polices, and ask if these converge to a limit. Or, one may ask if the associated affine shift $d_*(\eta)$ for the unconstrained process converges to a limit with vanishing η.

In case II*, with $\kappa = 1$, we find that $d_*(\eta)$ converges to the following identifiable limit:

$$d_* = \frac{1}{2}\frac{\sigma_H^2}{\delta_H} \log\left(1 + \frac{|\bar{c}_2^+|}{|\bar{c}_2^-|}\right), \tag{43}$$

where the height-process parameters are defined in (40).

Moreover, using the affine policy (37) with slope m_* and affine shift d_*, we find that the controlled CBM workload model is approximately pointwise-optimal *in the mean*. This is only possible in case II*: in case I the switching-curves are trivial, and independent of κ; in case III the associated height processes is not recurrent, so the construction of an optimal affine policy is not possible using the approach introduced here.

We first show that the lower switching-curve s_* is asymptotically affine, provided it is a smooth function of κ. A proof of proposition 4.5 is provided in the appendix.

Proposition 4.5. Assume that, for some $w_1^0 > 0$, the switching-curve $s_*(w_1^0; \cdot)$ is C^1 as a function of κ on [0, 1]. Let $d_0 \in \mathbb{R}$ denote the negative of the derivative at zero:

$$-d_0 := \lim_{\kappa \downarrow 0} \frac{1}{\kappa}\big[s_*\big(w_1^0; \kappa\big) - s_*\big(w_1^0; 0\big)\big].$$

Then the following limits hold:

(i) The optimal switching-curve is asymptotically affine:

$$\lim_{w_1 \to \infty} \big|s_*(w_1; 1) - \big(s_*(w_1; 0) - d_0\big)\big| = 0.$$

(ii) Suppose that $s_*(w_1^0; \cdot)$ admits a continuous second derivative, so that it is C^2 as a function of κ on [0, 1]. Then the rate of convergence may be quantified:

$$s_*(w_1; \kappa) = s_*(w_1, 0) - \kappa d_\circ + \mathrm{O}\big(\kappa^2\big), \quad \kappa \to 0, \text{ for each } w_1 > 0;$$
$$s_*(w_1; 1) = s_*(w_1, 0) - d_\circ + \mathrm{O}\big(w_1^{-1}\big), \quad w_1 \to \infty.$$

We show in theorem 4.7 that $d_* = d_\circ$ under the conditions of the proposition 4.5(ii). Moreover, in theorem 4.6 we show that this affine-shift parameter (approximately) minimizes $\mathsf{E}[\bar{c}(\widehat{W}(t))]$ over all affine policies, with quantifiable regret, uniformly on a fixed time interval of the form $t \in [T_1, T_2]$, with

$$0 < T_1 < T_2 < T_*(w) = \max_{i=1,2}\left(\frac{w_i}{1-\rho_i}\right).$$

Proofs of these results are contained in the appendix.

Theorem 4.6. Suppose that (A1)–(A3) are satisfied, and the process satisfies the conditions of case II*. Then, for any fixed $0 < T_1 < T_2 < T_*(w)$, there are constants $\varepsilon_0 > 0$, $b_0 < \infty$, $D_0 > 0$, such that for each $\|w\| = 1$, $|d - d_*| \leqslant 1$, we have approximate optimality

(i) For small κ: *for each* $\kappa \in (0, 1]$, $t \in [T_1, T_2]$,

$$\mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d_*)\big)\big] \leqslant \mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d)\big)\big] - \varepsilon_0\kappa|d-d_*|^2 + b_0 \exp\left(\frac{-D_0}{\kappa}\right)|d-d_*|.$$

(ii) For large initial conditions: *for each* $r \geqslant 1$, $t \in [rT_1, rT_2]$,

$$\mathsf{E}\big[\bar{c}\big(\widehat{W}(t; rw, 1, d_*)\big)\big] \leqslant \mathsf{E}\big[\bar{c}\big(\widehat{W}(t; rw, 1, d)\big)\big] - \varepsilon_0|d-d_*|^2 + b_0 \exp(-D_0 r)|d-d_*|.$$

Theorem 4.7. Suppose that (A1)–(A3) are satisfied, and the process satisfies the conditions of case II*. Assume in addition that, for some $w_1^0 > 0$,

$$s_*\big(w_1^0, \cdot\big) \text{ is } C^2 \text{ as a function of } \kappa \text{ on } [0, 1].$$

Then, there exists $b_0 < \infty$ independent of $w_1 > 0$ such that

$$\left|s_*(w_1; 1) - s(w_1; 1, d_*)\right| \leqslant b_0 \frac{1}{w_1}, \quad w_1 > 0.$$

That is, $d_* = d_\circ$, where $d_\circ$ is given in proposition 4.5.

5. Sensitivity

This section further develops structural properties of optimal control solutions for the CBM model. Our goal is to obtain either exact solutions or bounds on sensitivity with respect to the following variables,

- *Sensitivity with respect to variability.* We find that the optimal policy and cost scale linearly with κ.
- *Sensitivity with respect to policy.* A bounded perturbation to the optimal policy results in a vanishing "regret", of order κ^{-1}, as $\kappa \to \infty$.
- *Sensitivity to hard constraints.* Sensitivity with respect to buffer constraints is typically non-vanishing with increasing variability, and the regret may be unbounded as $\kappa \to \infty$.

Throughout this section we continue to assume that (A1)–(A3) hold. In particular, we restrict attention to two-dimensional workload models.

5.1. Sensitivity to variability

To understand the impact of variability in the model (19) we introduce a second parameterized family of network models by scaling the CBM workload process $\widehat{W}(t; w)$ with $\kappa = 1$. For any $\kappa > 0$ define

$$\widehat{W}^\kappa(t; w) = \kappa \widehat{W}\left(\frac{t}{\kappa}; \frac{w}{\kappa}\right).$$

This is also described as a linear model

$$\widehat{W}^\kappa(t; w) = w - (\mathbf{1} - \boldsymbol{\rho})t + I^\kappa(t) + N^\kappa(t), \quad \widehat{W}^\kappa(0) = w \in \mathsf{W}, \tag{44}$$

where $\boldsymbol{I}^\kappa$, $\boldsymbol{N}^\kappa$ are defined in the same way,

$$I^\kappa(t) := \kappa I\left(\frac{t}{\kappa}\right), \qquad N^\kappa(t) := \kappa N\left(\frac{t}{\kappa}\right), \quad t \in \mathbb{R}_+.$$

The following result is immediate from the definitions.

Lemma 5.1. For any $\kappa > 0$ the following two stochastic processes are identical in law:

$$N^\kappa(t) = \kappa N\left(\frac{t}{\kappa}\right), \qquad N(t; k) = \sqrt{\kappa} N(t), \quad t \geqslant 0.$$

Both are Brownian motion with zero drift, and instantaneous covariance equal to $\kappa \Sigma$.

This observation leads to an exact description of the dependency of optimal control laws on variability. Proposition 5.2 also provides a refinement of theorem 3.6. A proof is included in the appendix.

Proposition 5.2. Suppose that (A1)–(A3) hold for the model (19) (with $\kappa = 1$) so that a solution (h_*, γ_*) exists to the average cost optimality equations, and $\widehat{W}^*$ evolves in the set $\mathcal{R}_*$ shown in (35). Then, for any $\kappa > 0$ and any $w \in \mathsf{W}$,

(i) $\gamma_*(\kappa) = \kappa\gamma_*(1)$ is the optimal average cost for the model (19).

(ii) $h_*(w; \kappa) := \kappa^2 h_*(w/\kappa; 1)$ defines the relative value function for (19).

(iii) Assumption (A2) holds for the model (19), and the corresponding optimal region is given by,

$$\mathcal{R}_*(\kappa) := \kappa\mathcal{R}_* := \{\kappa w\colon\ w \in \mathcal{R}_*\}.$$

(iv) $s_*(w_1; \kappa) = \kappa s_*(w_1/\kappa; 1)$, $w_1 \in \mathbb{R}$, defines the optimal switching curve for (19).

(v) Suppose that $K \subset \mathcal{R}_*(0)$ is a given compact set, and that $h_*(w; \cdot)$ is C^2 on [0, 1] for each $w \in K$. Assume moreover that the derivatives with respect to κ are continuous on $K \times [0, 1]$.

Then, letting $\mathcal{R}_K$ denote the positive cone,

$$\mathcal{R}_K = \{aw\colon\ a \geqslant 0, w \in K\},$$

there exists a radially-homogeneous function $\ell\colon \mathcal{R}_K \to \mathbb{R}$ such that, for any $\kappa \geqslant 0$,

$$h_*(w, \kappa) = J_*(w) + \kappa\ell(w) + \mathrm{O}\big(\kappa^2\big), \quad w \in \mathcal{R}_K.$$

The error term $\mathrm{O}(\kappa^2)$ is bounded in w.

Note that we could alternatively consider a two-dimensional scaling to capture the impact of variability as well as load:

$$\widehat{W}(t; w, \kappa, \nu] := w - \nu(\mathbf{1} - \boldsymbol{\rho})t + I(t) + \sqrt{\kappa}N(t), \quad t \geqslant 0,$$

where (κ, ν) are non-negative scaling parameters. We then have

$$\nu^{-1}\widehat{W}(t; \nu w, \kappa, \nu) = w - (\mathbf{1} - \boldsymbol{\rho})t + \nu^{-1}I(t) + \nu^{-1}\sqrt{\kappa}N(t),$$

which is equal in distribution to $\widehat{W}(t; w, \nu^{-2}\kappa, 1)$ for a corresponding idleness policy. It follows that the optimal average cost satisfies,

$$\gamma_*(\kappa, \nu) = \frac{\kappa}{\nu}\gamma_*(1, 1), \quad \kappa > 0,\ \nu > 0.$$

Other properties of the relative value functions and optimal policies may also be established through this observation, as in the proof of proposition 5.2.

In the remainder of the paper we restrict attention to the model (19) in which load is fixed. In the next section we derive sensitivity formulae with respect to control parameters.

5.2. Sensitivity to policy structure

We begin with a qualitative analysis based on the fluid model $\widehat{w}$ for the example shown in figure 2. We consider the infinite-horizon control criterion, and an initial condition of the form $w^0 = (w_1^0, 0)^{\mathrm{T}}$ with $w_1^0 > 0$, and cost function shown in figure 3.

In case II* the sensitivity to policy structure may be significant. The optimal policy is the solution to an infinite-dimensional linear program, and consequently, sensitivity may be quantified through consideration of associated Lagrange multipliers. Even for smooth, nonlinear perturbations to s_* of the form considered in theorem 5.4 below, first-order sensitivity may be nonzero. This positive sensitivity *does not* hold in case III. The switching-curve s_* is an interior-point minimizer in case III, so that the first-order sensitivity is *zero* to any interior perturbation of the optimal policy.

For the CBM model one can again argue that first order sensitivity is zero. Moreover, the meantime spent near a boundary of $\mathcal{R}_*(\kappa)$ vanishes as $\kappa \to \infty$ or $\rho \to 1$. This is especially true in case III since the lower boundary of this set is 'repelling' (i.e. the height process is transient). In such cases one would expect to see far lower sensitivity of average cost with respect to changes in the switching-curves.

We now make these arguments precise through the following construction. Consider an arbitrary switching-curve that is *stabilizing* for the model when $\kappa = 1$: that is, we assume that the controlled process is positive recurrent, with finite steady-state cost $\gamma(1) \geqslant \gamma_*$. Define a new switching-curve for arbitrary κ by

$$s(w_1; \kappa) := \kappa s\left(\frac{w_1}{\kappa}\right), \quad w_1 \in \mathbb{R}_+.$$

The proof of the following result is identical to that of proposition 5.2.

Proposition 5.3. Suppose that (A1)–(A3) hold, and that for each $\kappa > 0$ the switching curve $s(\cdot; \kappa)$ is stabilizing for the model (19). Then, the steady state cost $\gamma(\kappa)$ is a linear function of κ:

$$\gamma(\kappa) = \kappa\gamma(1), \quad \kappa > 0.$$

Suppose now that a perturbation of this switching-curve is given. To understand its effect we again consider a parameterized family of policies. Suppose that $\Delta : \mathbb{R}_+ \to \mathbb{R}$ is a bounded, continuous function, and define for any real parameter $0 \leqslant \nu \leqslant 1$,

$$s(w_1; \kappa, \nu) := \kappa s\left(\frac{w_1}{\kappa}\right) + \nu\Delta\left(\frac{w_1}{\kappa}\right), \quad w_1 \in \mathbb{R}_+.$$

We have not multiplied Δ by κ since we wish to consider a perturbation of uniform size for the entire range of κ.

For example, suppose that s_* defines the optimal policy, and that the support of Δ is included in the support of s_*, in the sense that

$$\big(w_1, s_*(w_1; 1)\big) \notin \text{interior}(\mathsf{W}) \quad \Rightarrow \quad \Delta(w_1) = 0.$$

In this case, it follows from optimality that the steady-state cost $\gamma(\kappa, \nu)$ satisfies

$$\frac{\mathrm{d}}{\mathrm{d}\nu}\gamma(\kappa, 0) = 0 \quad \text{for all } \kappa > 0.$$

Theorem 5.4 shows that second order sensitivity vanishes as $\kappa \to \infty$. The proof again follows as in proposition 5.2.

Theorem 5.4. Suppose that (A1)–(A3) hold, and that the switching-curve $s(w_1; \kappa, \nu)$ is stabilizing when $\kappa = 1$, for any $0 \leqslant \nu \leqslant 1$, and that the steady-state cost $\gamma(1, \nu)$ is C^2 on [0, 1]. Then, the steady-state cost $\gamma(\kappa, \nu)$ satisfies

$$\gamma(\kappa, \nu) = \kappa\gamma\big(1, \nu\kappa^{-1}\big), \quad \kappa \geqslant 1,\ 0 < \nu \leqslant 1.$$

Consequently, for any $0 < \nu < 1$,

(i) Setting $\gamma_2 = (\mathrm{d}/\mathrm{d}\nu)\gamma(\kappa, \nu)$, we have

$$\begin{aligned} \big[\gamma(\kappa, \nu) - \gamma(\kappa, 0)\big] &\to \nu\gamma_2(1, 0), \\ \gamma_2(\kappa, \nu) &\to \gamma_2(1, 0), \quad \kappa \to \infty. \end{aligned} \tag{45}$$

(ii) Setting $\gamma_{22} = (\partial^2/\partial\nu^2)\gamma$,

$$\begin{aligned} &\gamma_{22}(\kappa, \nu) = \kappa^{-1}\gamma_{22}\big(1, \nu\kappa^{-1}\big), && \kappa \geqslant 1 \\ &\gamma_{22}(\kappa, \nu) \to 0, && \kappa \to \infty. \end{aligned}$$

5.3. Sensitivity to buffer constraints

Theorem 5.4 suggests that there is little reason to devote effort to exactly optimize a workload-relaxation since sensitivity is very low when κ is large. These theoretical results are plainly illustrated in the numerical results shown in section 6. This brings us back to the title of this paper – *where does the sensitivity lie?*

We restrict attention to the special case of buffer constraints. The main conclusion is that if certain buffer levels are constrained, independently of $\kappa > 0$, then the relative cost increases linearly with κ.

There is ample room in this area for further research. Structural results for the constrained control problem and some further numerical experiments are described in [12].

Consider first the fluid model. If buffer constraints are imposed then the state space takes the form $\mathsf{X} = \{x\colon\ x \geqslant \boldsymbol{\theta}, x \leqslant \boldsymbol{b}\}$. The workload space is equal to $\mathsf{W} = \{\widehat{\Xi}x\colon\ x \in \mathsf{X}\}$, and the effective cost $\bar{c}\colon \mathsf{W} \to \mathbb{R}_+$ is the solution to the linear program,

$$\begin{aligned} &\bar{c}(w) := \min r \\ \text{subject to}\quad & r \geqslant \langle c^i, x\rangle, \quad 1 \leqslant i \leqslant \ell_c, \\ & \Xi x = w, \\ & x \leqslant \boldsymbol{b}, \\ & x \geqslant \boldsymbol{\theta}. \end{aligned} \tag{46}$$

We again define $X^*(w)$ to be the optimizing $x \in \mathsf{X}$. The effective cost is a piecewise linear function in the variables $\{w_i, b_j\colon\ 1 \leqslant i \leqslant n,\ 1 \leqslant j \leqslant \ell\}$. This may be complex for a large network with many constraints, but it can be written explicitly [12].

Computing the sensitivity of cost with respect to a buffer constraint $b_i < \infty$ is straightforward given the formula (46):

$$\frac{\partial}{\partial b_i} c\big(\hat{q}^*(t; w)\big) = -\Gamma_i \mathbb{I}\big(\hat{q}_i^*(t; w) = b_i\big),$$

where $\Gamma_i \geqslant 0$ is a Lagrange-multiplier, and $\hat{\boldsymbol{q}}^*$ is the optimal solution to the relaxation on the constrained state space.

Consider now the stochastic workload-relaxation. Suppose that the policy is fixed, and that the state space W and the workload process $\widehat{\boldsymbol{W}}$ do not depend on $\boldsymbol{b}$ in a neighborhood of some value $\boldsymbol{b}_0$ of interest. In this case we obtain an exact expression for first-order sensitivity since only the effective cost c is subject to variation:

$$\frac{\partial}{\partial b_i} \mathsf{E}\big[\bar{c}\big(\widehat{W}(t)\big)\big] = -\Gamma_i \mathsf{P}\big(\widehat{W}(t) \in \mathsf{W}_{b_i}\big), \tag{47}$$

where $\mathsf{W}_{b_i} = \{w \in \mathsf{W}\colon\ \mathcal{X}^*(w)_i = b_i\}$. If the Lagrange multiplier is nonzero, and if the probability in (47) is non-vanishing for increasing κ, then we conclude that sensitivity with respect to this hard constraint is of order one as $\kappa \to \infty$. Moreover, under these assumptions we may conclude that for fixed $\boldsymbol{b}$,

$$\gamma^*(\kappa, \boldsymbol{b}) - \gamma^*(\kappa) = \kappa\left[\gamma^*\left(1, \frac{\boldsymbol{b}}{\kappa}\right) - \gamma^*(1)\right] \to \infty, \quad \kappa \to \infty.$$

Compare this with (45).

If the policy also changes with b_i then the formula (47) is no longer valid. However, theorem 5.4, and the numerical experiments presented in section 6 all suggest that the sensitivity of cost with respect to the policy in workload space is low when κ is large. Taking this for granted, the identity (47) suggests several approximations.

6. Numerical results

For a two-dimensional workload-relaxation we have seen that the optimized CBM model is a reflected Brownian motion on W in case I, for any $\kappa > 0$. For the original CRW queueing network model, an optimal solution will differ from the fluid solution by an approximately afnne perturbation of order $|\log(1-\rho)|$ (see [43, theorems 4.2, 4.3]). In cases II and III the stochastic switching-curve for the CRW model will differ substantially from the fluid solution: the gap is of order $|1-\rho|^{-1}$ (see the discussion following proposition 5.2).

To illustrate these conclusions, and other results from section 5, we compare optimal policies for several versions of the model shown in figure 2. Case I cannot hold since the effective cost is not monotone in this example. The two-dimensional relaxation for the fluid model may satisfy the conditions of case II or case III, depending upon the specific values of $\{\mu_i, \alpha\}$.

6.1. Models

We consider several sets of parameters $(\alpha, \mu_1, \mu_2, \mu_3)$, always with $\mu_1 = \mu_3$. For each set of parameters, we construct policies for each of three stochastic models:

Poisson model. The three-dimensional, countable state-space model (10) obtained through uniformization of a model with Poisson arrivals and exponential service distributions. The state process is denoted $\boldsymbol{Q}$.

Poisson workload-relaxation. The two-dimensional, countable state-space CRW model with state space $\mathsf{Y}_0 := \{y \in \mathbb{Z}_+^2 \colon\ y_2 \leqslant y_1\}$. This is described as the controlled random walk,

$$\widehat{Y}(k+1) = \widehat{Y}(k) - \boldsymbol{\delta} + I(k) + \widehat{N}(k), \quad k \in \mathbb{Z}_+,$$

where the drift vector is given by

$$-\boldsymbol{\delta} := \alpha \begin{pmatrix} 2 \\ 1 \end{pmatrix} - \mu_1 \begin{pmatrix} 1 \\ 0 \end{pmatrix} - \mu_2 \begin{pmatrix} 0 \\ 1 \end{pmatrix}; \tag{48}$$

and $\widehat{\boldsymbol{N}}$ is i.i.d. with zero mean, and common distribution given by

$$\widehat{N}(k) = \boldsymbol{\delta} + \begin{cases} 2e^1 + e^2 & \text{with probability } \alpha; \\ -e^1 & \text{with probability } \mu_1; \\ -e^2 & \text{with probability } \mu_2. \end{cases} \tag{49}$$

The two-dimensional idleness process $\boldsymbol{I}$ is adapted to $\widehat{\boldsymbol{Y}}$, and takes values in $\mathbb{Z}_+^2$.

Brownian workload-relaxation. The two-dimensional CBM model

$$\widehat{W}(t; w) = w - \boldsymbol{\delta} t + I(t) + N(t), \quad t \in \mathbb{R}_+, \tag{50}$$

whose parameters are defined consistently with $\widehat{\boldsymbol{Y}}$: the drift $\boldsymbol{\delta}$ is given in (48), and with $\widehat{\boldsymbol{N}}$ defined in (49),

$$\Sigma := t^{-1}\mathsf{E}\big[N(t)N(t)^{\mathrm{T}}\big] = \mathsf{E}\big[\widehat{N}(1)\widehat{N}(1)^{\mathrm{T}}\big]\alpha \begin{pmatrix} 4 & 2 \\ 2 & 1 \end{pmatrix} + \mu_1 \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} + \mu_2 \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} - \boldsymbol{\delta}\boldsymbol{\delta}^{\mathrm{T}}.$$

The Poisson models are chosen for purposes of computation since optimal policies can be obtained through value iteration. The Poisson workload-relaxation is a true relaxation of the scaled workload process for the Poisson model where

$$\begin{aligned} Y_1(k) &= \mu_1 W_1(k) = 2Q_1(k) + Q_2(k) + Q_3(k); \\ Y_2(k) &= \mu_2 W_2(k) = Q_1(k) + Q_2(k). \end{aligned} \tag{51}$$

Two fluid models are introduced for the purposes of policy synthesis:

Fluid model. The three-dimensional model (3), with state process $\boldsymbol{q}$, and system parameters given in (8).

Fluid workload-relaxation. The fluid model $\hat{\boldsymbol{y}}$ associated with the Poisson workload-relaxation is constrained to lie in the positive cone

$$\mathsf{Y} := \{y \in \mathbb{R}^2_+ \colon\ y_2 \leqslant y_1\}.$$

It satisfies the equations

$$\hat{y}_1(t) = \hat{y}_1(0) + (2\alpha - \mu_1)t + I_1(t); \qquad \hat{y}_2(t) = \hat{y}_2(0) + (\alpha - \mu_2)t + I_2(t), \quad (52)$$

where the associated two-dimensional cumulative-idleness process $\boldsymbol{I}$ has nondecreasing components.

The effective cost for the fluid relaxation takes a simple form when c is equal to the ℓ_1-norm. The following result is immediate from the definitions and proposition 3.2.

Proposition 6.1. Suppose that the cost function for the three-dimensional fluid model is given by $c(x) = x_1 + x_2 + x_3$, $x \in \mathsf{X}$. Then, the effective cost is equal to

$$\bar{c}(y) = \max(y_2, y_1 - y_2), \quad y \in \mathsf{Y}.$$

The associated monotone region is independent of the service and arrival rates, and is given by

$$\mathsf{Y}^+ := \left\{y \in \mathbb{R}^2_+ \colon\ \frac{1}{2}y_1 \leqslant y_2 \leqslant y_1\right\}.$$

In addition to presenting optimal policies in cases II and III, in this section we give results from a series of experiments to test sensitivity of cost with respect to control parameters. An optimal solution for the fluid workload-relaxation (52) is described by the *linear policy*,

$$\text{work resource 2 at maximal rate if } y_2 > m_* y_1,$$

for some constant $m_* \leqslant \frac{1}{2}$. For the Poisson workload-relaxation we consider affine policies of the specific form,

$$\text{work resource 2 at maximal rate if } (y_2 - \bar{y}_2) > m_*(y_1 - \bar{y}_1). \quad (53)$$

For the 3-dimensional Poisson model we consider various translations of (53) to buffer-coordinates.

6.2. *Numerical results for case II*

We consider two instances of case II for this model:

- CASE II(a). Rates: $\mu_1 = \mu_3 = 20$, $\mu_2 = 10$, $\alpha = 9$. The system is *balanced* with $\rho_1 = \rho_2 = \frac{9}{10}$.
- CASE II(b). Rates: $\mu_1 = \mu_3 = 20$, $\mu_2 = 11$, $\alpha = 9$. The loads are $\rho_1 = \frac{9}{10}$ and $\rho_2 = \frac{9}{11}$.

In each instance the vector $(1 - \rho_1, 1 - \rho_2)^{\mathrm{T}}$ lies *within* the monotone region W^+ (equivalently, the vector $(\mu_1 - 2\alpha, \mu_2 - \alpha)^{\mathrm{T}}$ lies within Y^+). Consequently, the optimal trajectory for the fluid workload-relaxation $\hat{\boldsymbol{y}}$ will be pathwise optimal. The optimal switching-curve for $\hat{\boldsymbol{y}}$ is defined by the switching-curve $y_2 = \frac{1}{2}y_1$ since this defines the lower boundary of Y^+.

We now turn to the CBM model (50). Case II(a) is a marginal case since the drift-vector $(2\alpha - \mu_1, \alpha - \mu_2)^{\mathrm{T}}$ is exactly aligned with the lower boundary of Y^+. Hence, the height process $\boldsymbol{H}$ associated with the CBM model is *not* positive recurrent in this case. Consequently, we would expect the optimal policy for the probabilistic workload-relaxations to differ significantly from the fluid relaxation (see (43)).

In contrast, in case II(b) we see that the conditions of case II* are met. Consideration of the height-process for the CBM model gives the following parameters in this case:

$$\delta_H = 1; \qquad \sigma_H^2 = 15,$$

and hence its steady-state distribution is exponential with parameter $\vartheta_H := 2\delta_H/\sigma_H^2 = 2/15$. The optimal affine shift (43) is given by

$$d_* = \frac{1}{2}15\log\big(1 + (1/1)\big) \approx 5.2.$$

In the numerical results below it will be more convenient to work with the corresponding w_1-axis intercept, given by

$$\beta_* := m_*^{-1} d_* = 2d_* \approx 10.4. \tag{54}$$

This is the asymptotic horizontal offset between the optimal switching-curves for the CBM model with $\kappa = 1$, and the fluid model, respectively.

Poisson workload-relaxation. Optimal policies for $\widehat{\boldsymbol{Y}}$ are shown in figure 1. Figure 1(a) shows the optimal policy for $\widehat{\boldsymbol{Y}}$ in case II(a) where the network is balanced. As expected, the difference between the fluid and stochastic switching curves is significant.

Figure 1(b) shows the optimal policy for $\widehat{\boldsymbol{Y}}$ in case II(b). It is closely approximated by an affine translation of the optimal policy for the fluid model. Moreover, the numerical results shown in figure 1(b) shows that the affine shift given in (54), which was obtained through consideration of the CBM model $\widehat{\boldsymbol{W}}$, results in an affine policy that coincides almost exactly with the the optimal policy for the Poisson model.

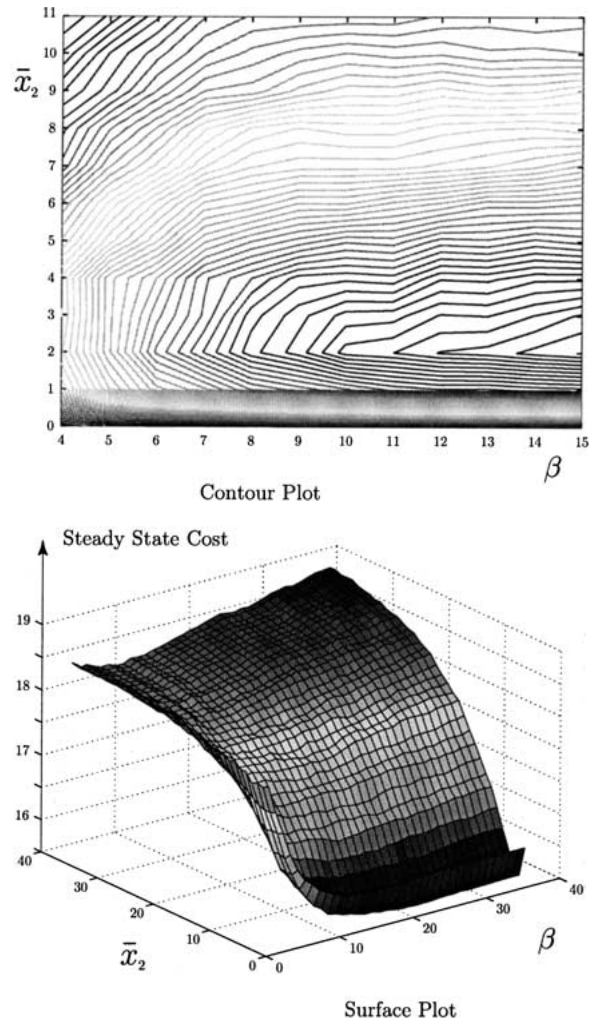

Figure 7. Contour and surface plots of average cost for case II(a).

Poisson model. We now compare a family of affine policies of the form (53) for the three-dimensional Poisson model. In case II these may be expressed:

$$\textit{Serve buffer one if } \text{buffer three is zero, or } x_3 \leqslant \beta \text{ and } x_2 \leqslant \bar{x}_2. \qquad (55)$$

For the fluid model, the optimal parameters are $\bar{x}_2 = 0$ and $\beta = \infty$.

Figure 7 shows results from simulation of the controlled network for a range of parameters in the affine policy (55) in case II(a). For each data-point the simulation was run for 5×10^6 time units, starting with an empty network. We see that the best affine shift β is very large. However, as may be predicted through theorem 5.4, sensitivity with respect to this parameter is extremely small.

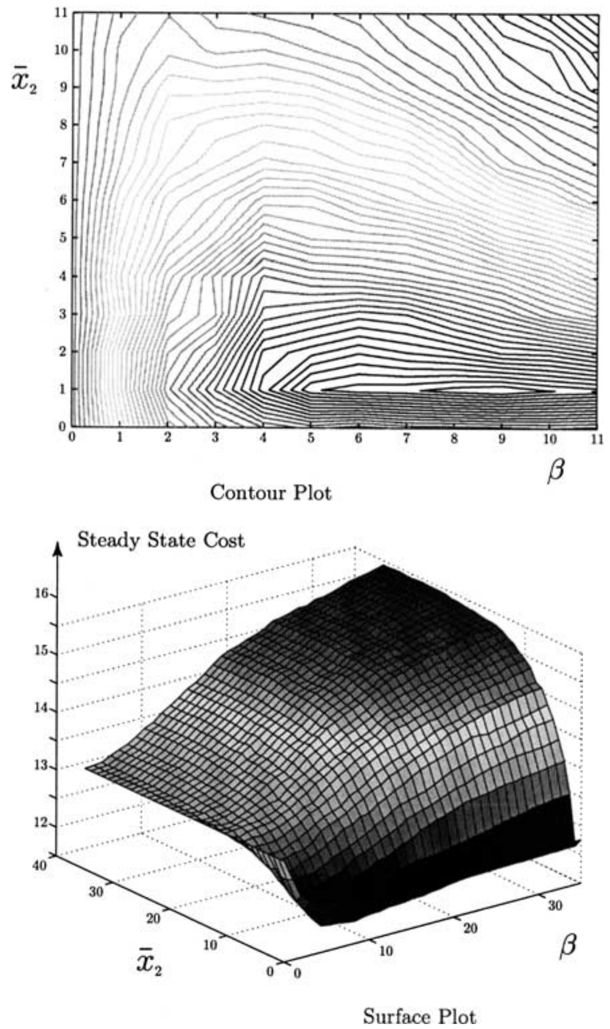

Figure 8. Contour and surface plots of average cost for case II(b).

Figure 8 shows the results from simulation experiments in case II(b). The optimal value of β is approximately equal to $\beta = 9$, which is consistent with the optimal offset for the Poisson workload-model $\widehat{\boldsymbol{Y}}$ shown in figure 1(b), and the asymptotic offset (54) for the optimized CBM model $\widehat{\boldsymbol{W}}$.

We can also define an affine policy for $\boldsymbol{Q}$ by explicitly translating an affine policy for $\widehat{\boldsymbol{Y}}$ as follows:

$$\textit{serve buffer one at time } k \textit{ if} \quad Y_2(k) \geqslant \frac{1}{2}(Y_1(k) - \beta),$$

where $\boldsymbol{Y}$ is given in (51). Equivalently, this policy is expressed,

$$\textit{serve buffer one at time } k \textit{ if} \quad Q_2(k) \geqslant Q_3(k) - \beta. \tag{56}$$

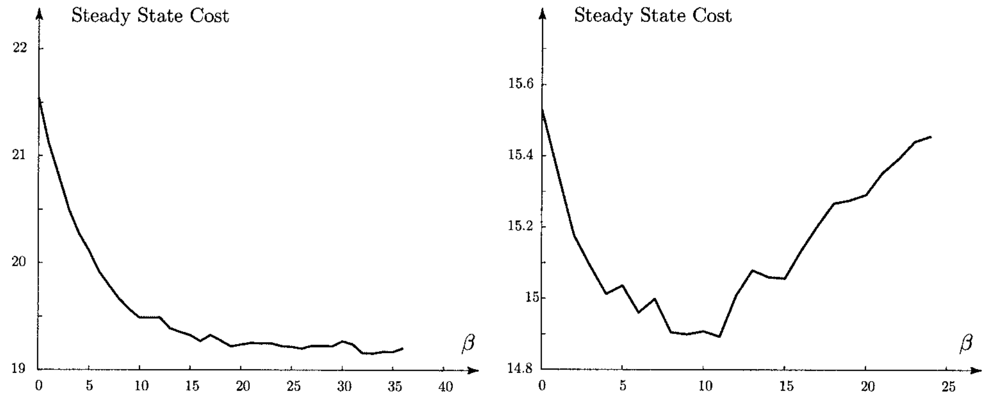

(a) Height process is null recurrent (b) Switching curve is attracting

Figure 9. Affine policies based on the Poisson workload-relaxation. The vertical axis shows the average cost obtained through a simulation of $\boldsymbol{Q}$ in case II for a time-horizon of $T = 5 \times 10^6$ samples.

Results from simulations of this policy are shown in figure 9. We see that the optimizing β is again consistent with theoretical results obtained for the CBM model.

The plots shown in figure 9 reveal significant variance even though the time-horizon was large: $T = 5 \times 10^6$ samples for each value of β. Variance in simulation of queues is typically high in moderate traffic [1,25–27,53]. The variance is more apparent in figure 9(b) due to the narrow range of values on the vertical axis.

6.3. *Numerical results for case III*

We now consider case III using the rates seen previously in (25). The loads at the two machines are $\rho_1 = 9/11$ and $\rho_2 = 9/10$, respectively. The theory in case III is far weaker than in the previous two examples where the conditions of case II are met. However, one would expect lower sensitivity with respect to the policy since the mean drift forces the process away from the optimal switching-curve for both stochastic and fluid models.

Although a pathwise optimal solution does not exist for the fluid model or its two-dimensional relaxation, the infinite-horizon optimal control for the fluid model is well known [2,51]. For this set of parameters it is expressed,

$$\textit{serve buffer one if } \text{buffer three is zero, or} \quad x_1 > x_3 \text{ and } x_2 = 0. \tag{57}$$

The optimal policy (57) enforces the following condition for $\boldsymbol{y}$:

$$\textit{work resource 2 at maximal rate if} \quad y_2 > \frac{1}{3} y_1. \tag{58}$$

An affine translation is implemented as follows:

$$\textit{serve buffer one if } \text{buffer three is zero, or} \quad 2x_2 + x_1 \geqslant x_3 - \beta \text{ and } x_2 \leqslant \bar{x}_2. \tag{59}$$

Numerical results in [39] show that (59) is a good fit with the optimal policy for $\boldsymbol{Q}$ when $\bar{x}_2 = 2$ and $\beta = 10$. Shown in figure 10 are the results of a simulation of these policies

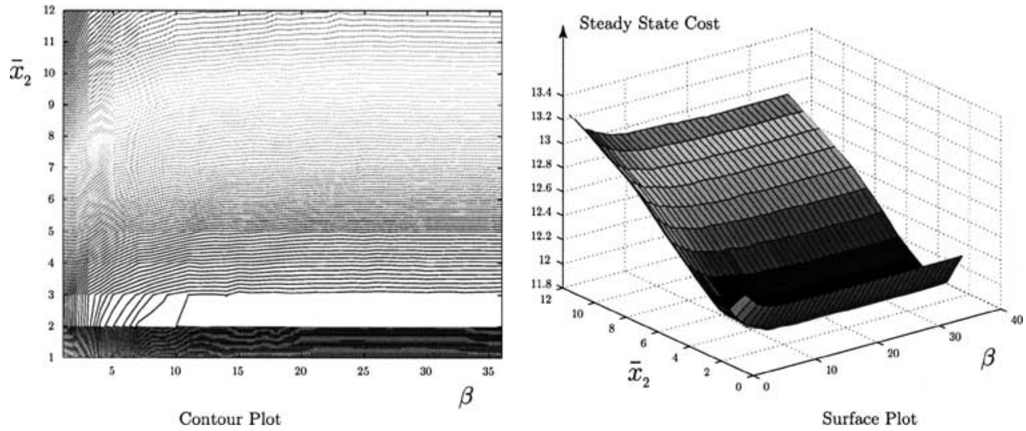

Figure 10. Contour and surface plots of average cost for various affine policies in case III.

with $\bar{x}_2 = 1, 2, \ldots, 12$ and $\beta = 1, 2, \ldots, 36$. The simulation was run for $T = 5 \times 10^6$ simulation steps.

This example illustrates the relative sensitivity of cost to hard constraints and interior points. On a fluid scale, the sensitivity of cost with respect to the threshold $x_2 \leqslant \bar{x}_2$ is strictly positive. On the other hand, for perturbations of the form considered in theorem 5.4, the first-order sensitivity is zero. The contour plot given in figure 10 shows that this dichotomy is inherited by the stochastic model in this example. Sensitivity with respect to β is *extremely low*, while sensitivity with respect to $\bar{x}_2$ is relatively high.

6.4. *Sensitivity to buffer constraints*

Given a vector $\boldsymbol{b} \geqslant \boldsymbol{\theta}$ of buffer constraints, the effective cost for the example under consideration is the solution to

$$
\begin{aligned}
\bar{c}(y) &= \min(x_1 + x_2 + x_3) \\
\text{subject to} \quad 2x_1 + x_2 + x_3 &= y_1, \\
x_1 + x_2 &= y_2, \\
x &\geqslant \boldsymbol{\theta}, \\
x &\leqslant \boldsymbol{b}.
\end{aligned}
$$

This has the explicit form

$$\bar{c}(y) = \max(y_2, y_1 - y_2, y_1 - b_1).$$

The effective cost is independent of b_2 and b_3, even though the workload space Y depends upon the values of these parameters.

We show below results from a single numerical experiment in case II(a). We take $b_2 = b_3 = \infty$ and vary the constraint b_1 on buffer one. The workload space is the same as in the unconstrained case, $\mathsf{Y} = \{y \in \mathbb{R}^2_+ \colon\ y_2 \leqslant y_1\}$.

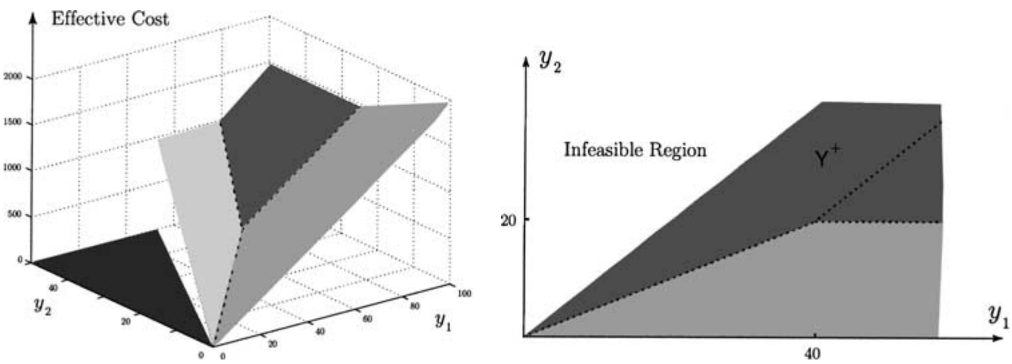

(a) (b)

Figure 11. (a) shows the effective cost for the three buffer model shown in figure 2 when the constraint $q_1 < b_1 = 20$ is imposed. The monotone region Y^+ is shown in (b). This set increases as the upper-bound b_1 decreases.

The effective cost and the monotone region Y^+ are shown in figure 11 in the case $b_1 = 20$. An optimal solution for the two-dimensional fluid relaxation will maintain $y^*(t; y) \in \mathsf{Y}^+$ for all $t > 0$, and $y \in \mathsf{Y}$. As b_1 decreases, the idle time at station 2 decreases in an optimal control solution. In the limiting case where $b_1 = 0$ the optimal solution $\hat{y}^*$ is pointwise minimal.

In the case of Poisson arrivals and exponential servers it is impossible to maintain a strict upper bound on Q_1 – we simply impose the constraint that buffer one receives priority whenever $Q_1(t) \geqslant b_1$. A policy of this form will maintain a constraint of the form $Q_1(t) \leqslant b_1 + K$ with high probability (of order $1 - (9/20)^K$). The control laws for $\boldsymbol{Q}$ considered in the numerical experiments below are modifications of those given in (55):

$$\textit{serve buffer one if} \quad \begin{array}{l} \text{buffer three is zero, or} \\ x_3 \leqslant \beta \quad \text{and} \quad x_2 \leqslant \bar{x}_2 \quad \text{and} \quad x_1 > 0, \quad \text{or} \\ x_1 \geqslant b_1. \end{array} \tag{60}$$

We stress that this is not a loss-model – we do not reject any arriving customers.

For this policy we use the expression (47) as a guide in an approximation of sensitivity with respect to b_1. Since $\Gamma_1 = 1$, a candidate sensitivity approximation is

$$\big[\mathsf{E}^{b_1=k}\big[c\big(Q(t)\big)\big] - \mathsf{E}^{b_1=k-1}\big[c\big(Q(t)\big)\big]\big] \approx -\big[\mathsf{P}^{b_1=k}\big(Q_1(t) \geqslant k\big)\big], \quad k \geqslant 1.$$

The right-hand side converges to zero geometrically fast as $k \to \infty$. For small k it is relatively large. The simplest case is $b_1 = 1$ since the resulting policy is First-Buffer First-Served (FBFS). In this case Q_1 is equivalent to an M/M/1 queue with arrival rate α and service rate μ_1. For the numerical values considered here it follows that

$$\mathsf{P}\big(Q_1(l) \geqslant 1\big) = \frac{\alpha}{\mu_1} = \frac{9}{20},$$

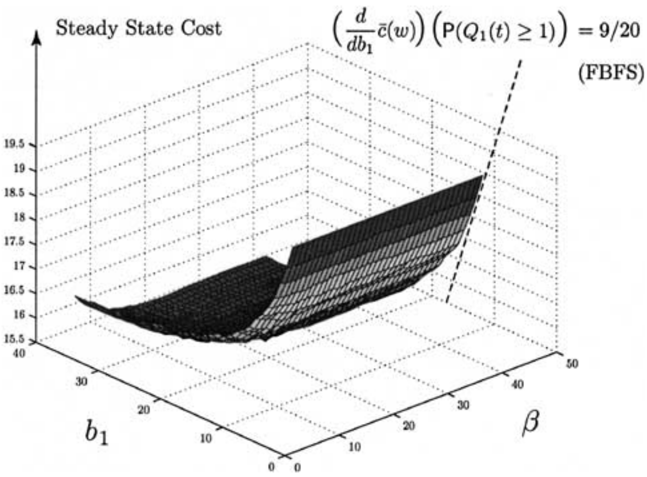

Figure 12. Surface plots of average cost for the Poisson model in case II(a) with buffer constraints. The sensitivity estimate (61) is nearly exact in this example. Sensitivity with respect to β remains very small.

and this gives the approximation,

$$\mathsf{E}^{b_1=1}\big[c\big(Q(t)\big)\big] - \mathsf{E}^{b_1=0}\big[c\big(Q(t)\big)\big] \approx -\frac{9}{20}. \tag{61}$$

The surface plot shown in figure 12 shows the steady-state mean $\mathsf{E}^{b_1=k}[c(Q(t))]$, computed via simulation, for a range of k under the policy (60). Again, for each datapoint in this plot, the simulation was run for 5×10^6 time units, starting with an empty network. We see in this example that the approximation (61) is nearly exact.

7. Conclusions

This paper has developed structural properties of various network models, and has broadened the solidarity between their respective optimal control solutions.

A central tool has been the analysis of a one-dimensional parameterized family of stochastic models. At one extreme, $\kappa = 0$, we obtain the fluid network-model, and as κ increases the model shows increasing variability. By exploring the relationship between these models we showed that increasing variability typically results in increasingly conservative optimal solutions (see, e.g., proposition 5.2). For example, for certain network models the fluid optimal solution would require significant idle time of bottlenecks in certain regimes (e.g., the sixteen-buffer model considered in [12,43]). With higher variability, an optimal solution will place higher priority on feeding bottlenecks. This corresponds exactly with the intuition of many working in the manufacturing area [19].

However, the conclusions of theorem 5.4 and the numerical results shown in section 6 all suggest that a focus on computing exact solutions to the workload optimization problem is not likely to yield significant improvements in applicability of the theory. In a specific application one will typically find control parameters for which sensitivity is far greater.

We are currently investigating in greater detail control issues surrounding buffer constraints; safety-stocks; and nonstandard performance metrics such as disaster-recovery [12,28]. These ideas may also be valuable in design. For example, *what is the true value of a reduction in variability?* The value of additional hardware or additional monitoring to reduce variability can be analyzed through a sensitivity analysis as described briefly in section 5.3.

We are also currently investigating a question posed in [43]: *Do the results of this paper lead to improved methods for performance approximation via simulation, or through calculation, by exploiting the simplicity of the network model following state space collapse?*

It is likely that a deeper look at the theory will lead to further insights. For example,

(i) The conclusions of theorem 4.6 may be extended to higher dimensions in the analog of 'case II*' where the resulting 'height-process' is positive recurrent. We do not know what happens in case III even in two dimensions, since in this case the switching surface is 'repelling'. We suspect that the qualitative form of the numerical results shown in section 6 will extend to higher dimensions since sensitivity to the switching surface is so low in this case.

(ii) A more detailed investigation of the associated differential generator for $\widehat{\boldsymbol{W}}$ is lacking in the present paper. This approach is pursued in [38] with impressively complex calculations. It may be possible to revisit this approach using the methods described here. For example, one can show that a fluid value function is in the domain of the generator for a given policy, thus providing a family of 'test functions' for analysis [6]. This may allow an extension of the model (19) to allow for the case where the variance Σ is a function of the current allocation rates.

(iii) If a realistic network model is constructed from renewal processes whose increments possess a moment generating function that is bounded in a neighborhood of the origin then a Gaussian approximation is typically justifiable in heavy-traffic, leading to a model of the form (19). When the driving noise does not satisfy a Central Limit Theorem then a Brownian motion approximation is no longer valid, and it may not be possible to construct a policy with finite steady-state cost. If the model includes a fractional Brownian motion $\boldsymbol{N}$ then higher levels of burstiness can be investigated, but one must move to a relaxed performance metric since steady-state cost is infinite when $\bar{c}$ is a norm.

Appendix A. Affine approximations

Proof of proposition 4.5. Under the C^1 assumption we have

$$s_*\big(w_1^0;\kappa\big) = s_*\big(w_1^0;0\big) - \beta_\circ\kappa + \mathcal{E}(\kappa), \quad 0 \leqslant \kappa \leqslant 1,$$

where the function $\mathcal{E}$ satisfies $\kappa^{-1}\mathcal{E}(\kappa) \to 0$, $\kappa \downarrow 0$. It then follows from proposition 5.2 that for arbitrary $m > 0$ and $0 \leqslant \kappa \leqslant m$,

$$\begin{aligned} s_*\big(mw_1^0; \kappa\big) &= ms_*\left(w_1^0; \frac{\kappa}{m}\right) \\ &= m\left\{s_*\big(w_1^0; 0\big) - d_\circ \frac{\kappa}{m} + \mathcal{E}\left(\frac{\kappa}{m}\right)\right\} \\ &= s_*\big(mw_1^0; 0\big) - d_\circ\kappa + \kappa\left(\frac{\kappa}{m}\right)^{-1}\mathcal{E}\left(\frac{\kappa}{m}\right). \end{aligned}$$

We thus obtain,

$$\left|s_*\big(mw_1^0; 1\big) - \big(s_*\big(mw_1^0; 0\big) - d_\circ\big)\right| = m\mathcal{E}\big(m^{-1}\big) \to 0, \quad m \to \infty,$$

and this immediately gives (i).

If s_* is C^2 in κ then $\mathcal{E}(\kappa) = \mathrm{O}(\kappa^2)$, which gives (ii). □

Recall from section 4 the definition of the controlled process $\widehat{W}(t; w^0, \kappa, d)$, using an affine policy with affine shift d and slope m_*. We let $X(t; w^0, \kappa, d)$ denote the unconstrained process using an affine policy with identical parameters.

In the results to follow we fix the initial condition w^0, assumed to satisfy $w_2^0 \leqslant m_* w_1^0 - \kappa d$ for all $d \in \mathbb{R}$ under consideration. The proof of theorem 4.6 is similar to theorem 4.3 in that we compare the constrained and unconstrained processes. Such comparisons are possible through bounds on the following first exit time,

$$\tau_d := \inf\big(t \geqslant 0\colon\ X\big(t; w^0, \kappa, d\big) \in \widehat{\mathsf{C}}^c\big).$$

The following result is immediate.

Lemma A.1. The random variable τ_d is a stopping time with respect to the filtration $\mathcal{F}_t = \sigma(N(s)\colon\ 0 \leqslant s \leqslant t)$, and

$$\widehat{W}\big(t; w^0, \kappa, d\big)\mathbb{I}_{\{t<\tau_d\}} = X\big(t; w^0, \kappa, d\big)\mathbb{I}_{\{t<\tau_d\}}.$$

Moreover, we have for all t, w^0, κ, d,

$$X\big(t; w^0, \kappa, d\big) = X\big(t; w^0, \kappa, 0\big) - (0, \kappa d)^{\mathrm{T}}.$$

In what follows we write $X(t; w^0, \kappa, 0)$ as $X(t)$ when it is not necessary to emphasize dependency on these parameters. We let $H(t)$ denote the associate height process, $H(t) = X_2(t) - m_* X_1(t)$, $t \geqslant 0$.

We begin with a result describing some statistical properties of the one-dimensional RBM $\boldsymbol{H}$:

Lemma A.2. Suppose that $\boldsymbol{H}$ is the reflected Brownian motion on $\mathbb{R}_+$ shown in (39), where $\boldsymbol{N}_H$ is a standard Brownian motion with instantaneous variance $\sigma_H^2 > 0, \sigma_H > 0$, and the reflection process $\boldsymbol{I}$ satisfies

$$\int_0^\infty H(t)\,\mathrm{d}I(t) = 0.$$

Then

(i) There exists $\varepsilon_0 > 0$ such that the process $\boldsymbol{H}$ is V-uniformly ergodic, with $V(x) := \mathrm{e}^{\varepsilon_0 x}$, $x \geqslant 0$ (see [16,44]).

(ii) The Markov process $\boldsymbol{H}$ possesses a stationary distribution π that is exponentially distributed with parameter $\vartheta_H := 2\delta_H/\sigma_H^2$.

(iii) The mean $\overline{H} = \int xp(x)\,\mathrm{d}x$ is given by $\overline{H} = \vartheta_H^{-1}$.

Proof. A proof of V-uniform ergodicity is included in [30]. The remaining results may be found in [11, theorem 6.2]. □

To bound the error between $\widehat{\boldsymbol{W}}$ and $\boldsymbol{X}$ on $[T_1, T_2]$ we apply the following lemma. The proof follows from consideration of an exponential martingale [48, section I.16].

Lemma A.3. Let $\boldsymbol{N}$ be a Brownian motion with zero drift and infinitesimal variance σ^2 and $N(0) = 0$. Then for any T_1, T_2 with $0 < T_1 < T_2$ and any $D > 0$ we have

$$\mathsf{P}\Big\{\sup_{t\in[T_1,T_2]} N(t) > D\Big\} \leqslant \exp\bigg(-\frac{1}{2\sigma^2}\bigg(\frac{D}{T_2 - T_1}\bigg)\bigg).$$

Proof of theorem 4.6. Part (ii) follows easily from part (i) and proposition 5.2, so only consider the case of vanishing κ.

Fix $d \in \mathbb{R}$, and define $\tau = \min(\tau_{d_*}, \tau_d)$, where d_* is defined in (43). We then write, for any $t \in [T_1, T_2]$,

$$\begin{aligned}
&\mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d_*)\big)\big] - \mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d)\big)\big] \\
&\quad = \mathsf{E}\big[\big(\bar{c}\big(\widehat{W}(t; w, \kappa, d_*)\big) - \bar{c}\big(\widehat{W}(t; w, \kappa, d)\big)\big)\mathbb{I}_{\{\tau<t\}}\big] \\
&\qquad - \mathsf{E}\big[\big(\bar{c}_X\big(X(t; w, \kappa, d_*)\big) - \bar{c}_X\big(X(t; w, \kappa, d)\big)\big)\mathbb{I}_{\{\tau<t\}}\big] \\
&\qquad + \mathsf{E}\big[\bar{c}_X\big(X(t; w, \kappa, d_*)\big) - \bar{c}_X\big(X(t; w, \kappa, d)\big)\big].
\end{aligned} \tag{A.1}$$

We will consider each of the three terms above separately.

To bound the first term we first note that since the function $\bar{c}$ is a norm on $\mathbb{R}^2$ and since all norms on $\mathbb{R}^2$ are equivalent, we must have, for some constant $b_0 > 0$,

$$\begin{aligned}
\big\|\bar{c}\big(\widehat{W}(t; w, \kappa, d_*)\big) - \bar{c}\big(\widehat{W}(t; w, \kappa, d)\big)\big\| &\leqslant b_0 \big\|\widehat{W}(t; w, \kappa, d_*) - \widehat{W}(t; w, \kappa, d)\big\| \\
&\leqslant b_0 \kappa \|d_* - d\|.
\end{aligned}$$

The first term is thus bounded by $b_0\kappa\|d_* - d\|\mathsf{P}[\tau < t]$.

To bound $\mathsf{P}[\tau < t]$, we apply lemma A.3. After transforming to the standard form required in the lemma we obtain, for some $D_1 > 0$ that depends on T_2,

$$\mathsf{P}[\tau < t] < \mathsf{P}[\tau < T_2] = \mathrm{O}\left(\exp\left(\frac{-D_1}{\kappa}\right)\right), \quad T_1 \leqslant t \leqslant T_2.$$

We conclude that the first term in (A.1) is bounded by $\|d_* - d\|\mathrm{O}(\exp(-D_1/\kappa))$. We may obtain an analogous bound on the second term where $\widehat{W}$ is replaced by X. The same arguments may be applied again since we have

$$\big\|X(t; w, \kappa, d) - X(t; w, \kappa, d_*)\big\| = \|d - d_*\|\kappa, \quad t > 0.$$

The analysis of the third term is a little more involved. Firstly, we use a Taylor's series expansion to write

$$\begin{aligned}
&\mathsf{E}\big[\bar{c}_X\big(X(t; w, \kappa, d)\big)\big] \\
&\quad = \mathsf{E}\big[\bar{c}_X\big(X(t; w, \kappa, d_*)\big)\big] + \left(\frac{\partial}{\partial d}\mathsf{E}\big[\bar{c}_X\big(X(t; w, \kappa, d_*)\big)\big]\right)(d - d_*) \\
&\qquad + \frac{1}{2}\left(\frac{\partial^2}{\partial d^2}\mathsf{E}\big[\bar{c}_X\big(X\big(t; w, \kappa, d^m\big)\big)\big]\right)(d - d_*)^2,
\end{aligned} \tag{A.2}$$

where d^m lies between d and d_*. The first derivative, evaluated at $d = d_*$, may be approximated as follows:

$$\begin{aligned}
\frac{\partial}{\partial d}\mathsf{E}\big[\bar{c}_X\big(X(t) - (0, \kappa d)^{\mathrm{T}}\big)\big] &= -\kappa\big\{c_2^-\mathsf{P}\big[H(t) < \kappa d\big] + c_2^+\mathsf{P}\big[H(t) \geqslant \kappa d\big]\big\} \\
&= -\kappa\big\{c_2^-\mathsf{P}[H < \kappa d] + c_2^+\mathsf{P}[H \geqslant \kappa d]\big\} + \mathrm{O}\left(\exp\left(\frac{-D_2 t}{\kappa}\right)\right)
\end{aligned} \tag{A.3}$$

for some $D_2 > 0$ independent of d_*. The last equation follows from V-uniform ergodicity of $\boldsymbol{H}$ (see lemma A.2). The random variable H has an exponential distribution with parameter $\vartheta_H = \vartheta_H(\kappa) = \vartheta_H(1)/\kappa$. This means that the probabilities $\mathsf{P}[H \geqslant \kappa d]$ and $\mathsf{P}[H < \kappa d]$ are independent of κ.

We also have, for a possibly smaller $D_3 > 0$,

$$\begin{aligned}
\frac{\partial^2}{\partial d^2}\mathsf{E}\big[\bar{c}_X\big(X(t) - (0, \kappa d)^{\mathrm{T}}\big)\big] &= -\kappa^2\big(c_2^- f_{H^1(t)}(\kappa d) - c_2^+ f_{H^1(t)}(\kappa d)\big) \\
&= -\kappa(c_2^- - c_2^+) f_{\kappa^{-1}H}(d) + \mathrm{O}\left(\exp\left(\frac{-D_3 t}{\kappa}\right)\right),
\end{aligned}$$

where the density $f_{\kappa^{-1}H}$ is independent of $\kappa > 0$. This shows that the second derivative in (A.2) is $\mathrm{O}(\kappa)$. It is also positive since $\bar{c}_X$ is convex.

Justification for differentiation of the expressions above follows from the assumption that $\Sigma > 0$. Consequently, for each $t \in [T_1, T_2]$, the random variable $H^1(t)$ has a strictly positive density function $f_{H^1(t)}(\cdot)$. It uniformly bounded in a neighborhood of zero for $t \in [T_1, T_2]$.

We now use the specific form of the parameter d_*. It is the solution to the equation

$$c_2^- \mathsf{P}[H < \kappa d_*] + c_2^+ \mathsf{P}[H \geqslant \kappa d_*] = 0.$$

It then follows from the previous bounds that there exists $b_1 < \infty$, $\varepsilon_1 > 0$ such that for all d in a neighborhood of d_*, and all $0 < \kappa \leqslant 1$,

$$\begin{aligned}&\mathsf{E}\big[\bar{c}_X\big(X\big(t; w^0, \kappa, d_*\big)\big)\big] - \mathsf{E}\big[\bar{c}_X\big(X\big(t; w^0, \kappa, d\big)\big)\big]\\&\quad\leqslant -\varepsilon_1 \kappa |d - d_*|^2 + b_1 \exp\left(\frac{-D_2 t}{\kappa}\right)\kappa |d - d_*|.\end{aligned}$$

Combining this with the bounds obtained for the three terms on the right-hand side of (A.1) one arrives at the penultimate bound: for some $D_2 > 0$, $\varepsilon_2 > 0$, $b_2 < \infty$, and all $0 \leqslant \kappa \leqslant 1$, $T_1 \leqslant t \leqslant T_2$,

$$\begin{aligned}&\mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d_*)\big)\big] - \mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w, \kappa, d)\big)\big]\\&\quad\leqslant -\varepsilon_2 \kappa |d - d_*|^2 + b_2 \exp\left(\frac{-D_2}{\kappa}\right)|d - d_*|.\end{aligned} \tag{A.4}$$

This gives the desired bound in theorem 4.6. □

Proof of theorem 4.7. An application of proposition 4.5 ensures that the desired bound holds

$$\big|s_*(w_1; 1) - s(w_1; 1, d_\circ)\big| \leqslant b_0 \frac{1}{w_1}, \quad w_1 > 0. \tag{A.5}$$

To prove the result we demonstrate that $d_\circ = d_*$.

Consider a large initial condition rw, where $r \geqslant 1$, and $w = (1, m)^{\mathrm{T}} \in \mathsf{W}$ with $m < m_*$. Hence, under the affine policy with parameter d, we have $\widehat{W}(0+; rw, d) = (r, m_* r - d)^{\mathrm{T}}$ for all $r \geqslant 0$ sufficiently large.

We fix $\varepsilon > 0$ and set $T_2 = T_*(w) - \varepsilon$. In the case where $d = d_\circ$, we apply (A.5) to obtain the approximation,

$$\begin{aligned}h_*(rw) &= \mathsf{E}\left[\int_0^{rT_2} \big[\bar{c}\big(\widehat{W}^*(s; rw)\big) - \gamma_*\big]\,\mathrm{d}s + h_*\big(\widehat{W}^*(rT_2; rw)\big)\right]\\&= \mathsf{E}\left[\int_0^{rT_2} \big[\bar{c}\big(\widehat{W}(s; rw, d_\circ)\big) - \gamma_*\big]\,\mathrm{d}s + h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\right]\\&\quad + \mathrm{O}(1) + \mathsf{E}\big[h_*\big(\widehat{W}^*(rT_2; rw)\big) - h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\big].\end{aligned}$$

Combining (A.5) and theorem 3.6(iii) with this bound then gives

$$h_*(rw) = \mathsf{E}\left[\int_0^{rT_2} \big[\bar{c}\big(\widehat{W}(s; rw, d_\circ)\big) - \gamma_*\big]\,\mathrm{d}s + h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\right] + \mathrm{O}(1). \tag{A.6}$$

By optimality we also have, for *any* $d \in \mathbb{R}^2_+$,

$$h_*(rw) = \mathsf{E}\left[\int_0^{rT_2} \left[\bar{c}\big(\widehat{W}(s; rw, d)\big) - \gamma_*\right] \mathrm{d}s + h_*\big(\widehat{W}(rT_2; rw, d)\big)\right].$$

Specializing to $d = d_*$, and applying theorem 4.6 together with the same arguments used above then gives

$$\begin{aligned}
h_*(rw) &\leqslant \mathsf{E}\left[\int_0^{rT_2} \left[\bar{c}\big(\widehat{W}(s; rw, d_\circ)\big) - \gamma_*\right] \mathrm{d}s + h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\right] \\
&\quad - \mathrm{O}(r)\frac{\|d_\circ - d_*\|^2}{1 + \|d_\circ - d_*\|} + \mathrm{O}(1) \\
&\quad + \mathsf{E}\left[h_*\big(\widehat{W}^*(rT_2; rw, d_*)\big) - h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\right] \\
&\leqslant \mathsf{E}\left[\int_0^{rT_2} \left[\bar{c}\big(\widehat{W}(s; rw, d_\circ)\big) - \gamma_*\right] \mathrm{d}s + h_*\big(\widehat{W}(rT_2; rw, d_\circ)\big)\right] \\
&\quad - \mathrm{O}(r)\frac{\|d_\circ - d_*\|^2}{1 + \|d_\circ - d_*\|} + \mathrm{O}(r)\|d_\circ - d_*\|\varepsilon + \mathrm{O}(1).
\end{aligned}$$

The negative term appears due to optimality of d_*.

Consequently, on combining this bounds with (A.6), we have

$$h_*(rw) \leqslant h_*(rw) - \mathrm{O}(r)\frac{\|d_\circ - d_*\|^2}{1 + \|d_\circ - d_*\|} + \mathrm{O}(r)\|d_\circ - d_*\|\varepsilon + \mathrm{O}(1).$$

This proves the result that $d_\circ = d_*$ since $\varepsilon > 0$ is arbitrary. □

Appendix B. A proof of proposition 5.2

All of these results are based on lemma 5.1, which justifies the consideration of (44) in place of (19).

For any $\kappa > 0$, the optimal average cost $\widehat{K}_*(w; \kappa)$ is independent of the initial condition $w \in \mathsf{W}$. Also, on considering the model (44), a change of variables $u = t\kappa$ then shows

$$\begin{aligned}
\gamma_* &= \inf\left(\lim_{T\to\infty} \mathsf{E}\left[\frac{1}{T}\int_0^T \bar{c}\left(\widehat{W}\left(t; \frac{w}{\kappa}\right)\right) \mathrm{d}t\right]\right) \\
&= \kappa^{-1}\left\{\inf\left(\lim_{T\to\infty} \mathsf{E}\left[\frac{1}{\kappa T}\int_0^{\kappa T} \bar{c}\big(\widehat{W}^\kappa(u; w)\big)\,\mathrm{d}u\right]\right)\right\},
\end{aligned}$$

where the infimum is over all admissible policies. This gives (i), and the proof of (ii) is identical.

To see (iii) write $\gamma_* = \gamma_*(1)$, and apply (i) to the definition (33),

$$
\begin{aligned}
h_*(w) &:= \inf \int_0^\infty \mathsf{E}\big[\bar{c}\big(\widehat{W}(t; w)\big) - \gamma_*\big]\,\mathrm{d}t \\
&= \kappa^{-1}\left(\inf \int_0^\infty \mathsf{E}\left[\bar{c}\left(\widehat{W}\left(\frac{u}{\kappa}; w\right)\right) - \frac{\gamma_*(\kappa)}{\kappa}\right]\mathrm{d}u\right) \\
&= \kappa^{-2}\left(\inf \int_0^\infty \mathsf{E}\big[\bar{c}\big(\widehat{W}^\kappa(u; \kappa w)\big) - \gamma_*(\kappa)\big]\,\mathrm{d}u\right) \\
&= \kappa^{-2} h_*(\kappa w; \kappa),
\end{aligned}
$$

which is (iii).

The proof of (iv) then follows from the previous identities. A second-order Taylor-series expansion gives for any $w \in K$,

$$
h_*(w; \kappa) = h_*(w; 0) + \kappa \ell(w) + \frac{1}{2}\frac{\partial^2}{\partial \kappa^2} h_*(w; \kappa_0)\kappa^2,
$$

where $0 < \kappa_0(w) < \kappa$, and $\ell(w) = (\partial/\partial\kappa) h_*(w; 0)$. Radial-homogeneity of the function ℓ on $\mathcal{R}_K$ follows from (iii): for any $a > 0$, $w \in K$,

$$
\frac{\partial}{\partial \kappa} h_*(aw; \kappa) = \frac{\partial}{\partial \kappa}\left[a^2 h_*\left(w; \frac{\kappa}{a}\right)\right] = a \frac{\partial}{\partial \kappa} h_*\left(w; \frac{\kappa}{a}\right).
$$

Setting $\kappa = 0$ then gives $\ell(aw) = a\ell(w)$.

Since K is compact we have a uniform bound

$$
h_*(w; \kappa) = \hat{J}_*(w) + \kappa \ell(w) + \mathrm{O}\big(\kappa^2\big), \quad w \in K.
$$

For arbitrary nonzero $aw \in \mathcal{R}_K$, with $w \in K$, we then have from (iii)

$$
\begin{aligned}
h_*(aw, \kappa) &= a^2 h_*\left(w, \frac{\kappa}{a}\right) \\
&= a^2\left[\hat{J}_*(w) + \frac{\kappa}{a}\ell(w) + \mathrm{O}\left(\left(\frac{\kappa}{a}\right)^2\right)\right] \\
&= \hat{J}_*(aw) + \kappa \ell(aw) + \mathrm{O}\big(\kappa^2\big).
\end{aligned}
$$

References

[1] S. Asmussen, Queueing simulation in heavy traffic, Math. Oper. Res. 17 (1992) 84–111.

[2] F. Avram, D. Bertsimas and M. Ricard, An optimal control approach to optimization of multiclass queueing networks, in: *Proc. of Workshop on Queueing Networks of the Mathematical Institute*, eds. F. Kelly and R. Williams, Minneapolis, 1994, IMA Volumes in Mathematics and its Applications, Vol. 71 (Springer, New York, 1995).

[3] F.L. Baccelli, G. Cohen and G.J. Olsder, *Synchronization and Linearity: An Algebra for Discrete Event Systems*, Wiley Series in Probability and Mathematical Statistics (Wiley, New York, 1992).

[4] S.L. Bell and R.J. Williams, Dynamic scheduling of a system with two parallel servers: Asymptotic optimality of a continuous review threshold policy in heavy traffic, in: *Proc. of the 38th Conf. on Decision and Control*, Phoenix, AZ, 1999, pp. 1743–1748.

[5] P. Billingsley, *Convergence of probability measures*, Wiley Series in Probability and Statistics: Probability and Statistics, 2nd ed. (Wiley-Interscience, New York, 1999).
[6] V.S. Borkar and S.P. Meyn, Value functions and simulation in stochastic networks, in: *42th IEEE Conf. on Decision and Control*, 2003, submitted.
[7] M. Bramson, State space collapse with application to heavy traffic limits for multiclass queueing networks, Queueing Systems 30 (1998) 89–148.
[8] M. Bramson and R.J. Williams, On dynamic scheduling of stochastic networks in heavy traffic and some new results for the workload process, in: *Proc. of the 39th Conf. on Decision and Control*, 2000.
[9] H. Chen and A. Mandelbaum, Discrete flow networks: bottleneck analysis and fluid approximations, Math. Oper. Res. 16(2) (1991) 408–446.
[10] H. Chen and D.D. Yao, Dynamic scheduling of a multiclass fluid network, Oper. Res. 41(6) (1993) 1104–1115.
[11] H. Chen and D.D. Yao, *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*, Stochastic Modelling and Applied Probability (Springer, New York, 2001).
[12] M. Chen, R. Dubrawski and S.P. Meyn, Management of demand-driven production systems (2002) submitted for publication.
[13] R.L. Cruz, A calculus for network delay, part I: Network elements in isolation, IEEE Trans. Inform. Theory 31 (1991) 114–131.
[14] J.G. Dai, On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models, Ann. Appl. Probab. 5(1) (1995) 49–77.
[15] J.G. Dai and S.P. Meyn, Stability and convergence of moments for multiclass queueing networks via fluid limit models, IEEE Trans. Automat. Control 40 (November 1995) 1889–1904.
[16] D. Down, S.P. Meyn and R.L. Tweedie, Exponential and uniform ergodicity of Markov processes, Ann. Probab. 23(4) (1995) 1671–1691.
[17] P. Dupuis and H. Kushner, *Numerical Methods for Stochastic Control Problems in Continuous Time*, Applications of Mathematics, Vol. 24 (Springer, New York, 2001).
[18] J.D. Eng, Humphrey and S.P. Meyn, Fluid network models: Linear programs for control and performance bounds, in: *Proc. of the 13th IFAC World Congress*, eds. J. Cruz, J. Gertler and M. Peshkin, Vol. B, San Francisco, CA, 1996, pp. 19–24.
[19] S.B. Gershwin, *Manufacturing Systems Engineering* (Prentice-Hall, Englewood Cliffs, NJ, 1993).
[20] J.M. Harrison, *Brownian Motion and Stochastic Flow Systems* (Wiley, New York, 1985).
[21] J.M. Harrison, Brownian models of queueing networks with heterogeneous customer populations, in: *Stochastic Differential Systems, Stochastic Control Theory and Applications*, Minneapolis, MN, 1986 (Springer, New York, 1988) pp. 147–186.
[22] J.M. Harrison, Brownian models of open processing networks: Canonical representations of workload, Ann. Appl. Probab. 10 (2000) 75–103.
[23] J.M. Harrison, Stochastic networks and activity analysis, in: *Analytic Methods in Applied Probability, In Memory of Fridrih Karpelevich*, ed. Y. Suhov (Amer. Math. Soc., Providence, RI, 2002).
[24] J.M. Harrison and L.M. Wein, Scheduling networks of queues: Heavy traffic analysis of a two-station closed network, Oper. Res. 38(6) (1990) 1052–1064.
[25] S.G. Henderson, Variance reduction via an approximating Markov process, Ph.D. thesis, Stanford University, Stanford, CA, USA (1997).
[26] S.G. Henderson and P.W. Glynn, Approximating martingales for variance reduction in Markov process simulation, Math. Oper. Res. 27(2) (2002) 253–271.
[27] S.G. Henderson and S.P. Meyn, Variance reduction for simulation in multiclass queueing networks, IIE Trans. Oper. Engrg. (2003) to appear.
[28] S.G. Henderson, S.P. Meyn and V. Tadic, Performance evaluation and policy selection in multiclass networks, *Discrete Event Dynamic Systems: Theory and Applications*, Special Issue on Learning and Optimization Methods in Discrete Event Dynamic Systems (2002) to appear.

[29] F.C. Kelly and C.N. Laws, Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling, Queueing Systems 13 (1993) 47–86.

[30] I. Kontoyiannis and S.P. Meyn, Spectral theory and limit theorems for geometrically ergodic Markov processes, *Ann. Appl. Probab.* 13 (2003) 304–362; presented at *The INFORMS Applied Probability Conference*, New York, July 2001.

[31] S. Kumar and M. Muthuraman, A numerical method for solving singular Brownian control problems, in: *Proc. of the 39th Conf. on Decision and Control*, 2000.

[32] H.J. Kushner, *Heavy Traffic Analysis of Controlled Queueing and Communication Networks*, Stochastic Modelling and Applied Probability (Springer, New York, 2001).

[33] H.J. Kushner and K.M. Ramchandran, Optimal and approximately optimal control policies for queues in heavy traffic, SIAM J. Control Optim. 27 (1989) 1293–1318.

[34] X. Luo and D. Bertsimas, A new algorithm for state-constrained separated continuous linear programs, SIAM J. Control Optim. 37 (1998) 177–210.

[35] C. Maglaras, Dynamic scheduling in multiclass queueing networks: Stability under discrete-review policies, Queueing Systems 31 (1999) 171–206.

[36] C. Maglaras, Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimally, Ann. Appl. Probab. 10 (2000).

[37] L.F. Martins and H.J. Kushner, Routing and singular control for queueing networks in heavy traffic, SIAM J. Control Optim. 28 (1990) 1209–1233.

[38] L.F. Martins, S.E. Shreve and H.M. Soner, Heavy traffice convergence of a controlled, multiclass queueing system, SIAM J. Control Optim. 34(6) (1996) 2133–2171.

[39] S.P. Meyn, The policy iteration algorithm for average reward Markov decision processes with general state space, IEEE Trans. Automat. Control 42 (1997); also presented at *The 35th IEEE Conf. on Decision and Control*, Kobe, Japan, December 1996.

[40] S.P. Meyn, Stability and optimization of queueing networks and their fluid models, in: *Mathematics of Stochastic Manufacturing Systems*, Williamsburg, VA, 1996 (Amer. Math. Soc., Providence, RI, 1997) pp. 175–199.

[41] S.P. Meyn, Sequencing and routing in multiclass queueing networks. Part I: Feedback regulation, SIAM J. Control Optim. 40(3) (2001) 741–776.

[42] S.P. Meyn, Stability, performance evaluation, and optimization, in: *Markov Decision Processes: Models, Methods, Directions, and Open Problems*, eds. E. Feinberg and A. Shwartz (Kluwer, Dordrecht, 2001) pp. 43–82.

[43] S.P. Meyn, Sequencing and routing in multiclass queueing networks. Part II: Workload relaxations, SIAM J. Control Optim. (2003) to appear; also presented at *The 2000 IEEE Internat. Symposium on Information Theory*, Sorrento, Italy, June 25–June 30 2003.

[44] S.P. Meyn and R.L. Tweedie, *Markov Chains and Stochastic Stability* (Springer, London, 1993).

[45] J. Perkins, Control of push and pull manufacturing systems, Ph.D. thesis, University of Illinois, Urbana, IL (1993), Technical Report No. UILU-ENG-93-2237 (DC-155).

[46] M.I. Reiman, Open queueing networks in heavy traffic, Math. Oper. Res. 9 (1984) 441–458.

[47] M.I. Reiman, A multiclass queue in heavy traffic, Adv. in Appl. Probab. 20 (1988) 179–207.

[48] L.C.G. Rogers and D. Williams, *Diffusions, Markov Processes, and Martingales*, Vol. I: *Foundations*, 2nd ed. (Cambridge Univ. Press, Cambridge, 2000).

[49] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis: Queues, Communication and Computing* (Chapman and Hall, London, UK, 1995).

[50] M.H. Veatch, Using fluid solutions in dynamic scheduling (2001) submitted for publication.

[51] G. Weiss, Optimal draining of a fluid re-entrant line, in: *IMA Volumes in Mathematics and its Applications*, Vol. 71, eds. F. Kelly and R. Williams (Springer, New York, 1995) pp. 91–103.

[52] G. Weiss, A simplex based algorithm to solve separated continuous linear programs, Technical Report, Department of Statistics, University of Haifa, Israel (2001).

[53] W. Whitt, Planning queueing simulations, Managm. Sci. 35 (1994) 1341–1366.