In search of sensitivity in network optimization*

Mike Chen, Charuhas Pandit, and Sean Meyn†

June 13, 2002

Abstract

This paper concerns policy synthesis in large queuing networks. The results provide answers to the following questions:

(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 use 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? If so, what is the 'regret'?

(iii) Where are the highest sources of 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. However, sensitivity of steady state performance with respect to perturbations in the policy vanishes with increasing variability.

Keywords
Queuing Networks, Routing, Scheduling, Optimal Control.

AMS subject classifications
Primary: 90B35, 68M20, 90B15.
Secondary: 93E20.

*Work supported in part by grants DMI 00 85165 and ECS 99 72957

†Department of Electrical and Computer Engineering and the Coordinated Sciences Laboratory, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A. (s-meyn@uiuc.edu).

1 Introduction

Optimization and performance evaluation of networks is of great interest in both academia and the industry for obvious reasons. It is equally obvious that standard approaches to optimization based upon a discrete state-space Markov queueing model lead to intractable optimality equations even for small network models. This has led to the development of various alternative network models tuned to address particular issues such as steady-state performance; contention for resources; or the impact of breakdowns (e.g. [14, 15, 6, 2, 37, 7].)

The simplest model is the linear, deterministic fluid model used in, for example, [1, 5, 4, 8, 24, 25, 26, 34, 39, 38]. It provides a framework for policy synthesis for large networks based on linear programming, and this leads to attractive approaches to sensitivity analysis through the associated Lagrange-multipliers.

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 [9, 10]. Results establishing solidarity among respective optimal control solutions are developed in [30, 31], based on this stability theory. It is shown that scaled optimal solutions for the stochastic network are approximated by the optimal solution for the fluid model. Conversely, it is shown in [32] that 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.

The linear 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, [35, 15, 36, 16, 23, 27, 19, 15, 20, 6, 22].

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 [19], and further studied in several subsequent references. When the effective cost is monotone then one obtains a policy that is approximately path-wise optimal in heavy traffic [28]. A similar approach is pursued in [3].

When the effective cost for the fluid-model is not monotone then an optimal policy is *not* path-wise optimal for the Brownian model (see [32, Section 4.5]). An optimal policy for the Brownian model is defined by nonlinear switching curves in workload-space. In this case only qualitative structural results have been established in small examples (e.g. [28]), and numerical studies have appeared in [21, 12].

The aforementioned optimality theory is based upon a workload representation of the network under study. Related general constructions are described in [17, 31], and this framework will form the basis of the results of the present paper. Some of the issues to be addressed are listed below.

(i) How does a policy for a stochastic model change when variability is increased? When the monotonic effective cost assumption of [31] does not hold we demonstrate in Proposition 3.3 that an optimal policy scales linearly with increasing variability.

(ii) When the effective cost is not monotone, but the fluid model admits a path-wise optimal solution, we construct an *affine policy* based on the optimal fluid solution that is approximately path-wise optimal in the mean, with exponentially small regret (see Theorem 4.2).

(iii) In Theorem 4.4 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 correspond to hard constraints in the deterministic optimal control problem. Some results are described in Section 4.3, and illustrated in Section 5.

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, and this section also contains new results that strengthen the solidarity first established in [30]. In the development of Section 4 we establish qualitative bounds on sensitivity with respect to control parameters and with respect to system variability. Section 5 provides detailed numerical examples, and Section 6 contains conclusions and suggestions for further research.

2 Network models and their relaxations

We begin with a description of the basic network models.

2.1 Network models in 'buffer-coordinates'

The results of this paper are based on two primary network models: A stochastic network model and its fluid counterpart.

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 fluid model are tremendous virtues in control design.

The stochastic and fluid network models are described through the respective equations,

$$Q(t;x) = x - S(Z(t;x)) + R(Z(t;x)) + A(t) \tag{1}$$

$$q(t;x) = x + Bz(t;x) + \alpha t, \qquad t \in \mathbb{R}_+ \tag{2}$$

In both models time is continuous; the *state processes* $Q(t;x)$, $q(t;x)$ evolve on $\mathsf{X} := \mathbb{R}^{\ell}_+$; and the *allocation processes* $Z(t;x)$, $z(t;x)$ evolve on $\mathbb{R}^{\ell_u}_+$, for some integers $\ell, \ell_u \geq 1$.

For the fluid model (2) we have the following interpretations,

(i) $q(t;x) \in \mathbb{R}^{\ell}_+$ is a vector of buffer-levels of various materials in the network;

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

$$\zeta(t;x) \in \mathsf{U}, \qquad \mathsf{U} := \{u \in \mathbb{R}^{\ell_u}_+ : Cu \leq \mathbf{1}\} \qquad t \in \mathbb{R}_+, \tag{3}$$

where the *constituency matrix* C is an $\ell_m \times \ell_u$ matrix with binary entries, and $\mathbf{1}$ denotes a vector of ones.

(iii) The vector α 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 $B = -S + R$ where R, S have positive entries: The value S_{ij} is the rate at which buffer i is drained when activity j is permitted to work at maximum rate $\zeta_j(t;x) = 1$, and R_{kj} is the rate at which that material is then sent to buffer k.

The stochastic network model (1) is a version of the stochastic processing network developed in [17, 18]. It is subject to analogous interpretations and constraints:

(i) $Q(t;x) \in \mathbb{R}_+^\ell$ is again a vector of buffer-levels. Typically, the entries of $\boldsymbol{Q}$ are further constrained to be integer-valued, and the entries of $\boldsymbol{Z}$ are piece-wise linear, but these assumptions are not required in this paper.

(ii) The allocation process $\boldsymbol{Z}$ satisfies the linear constraints,

$$\frac{Z(t;x) - Z(s;x)}{t-s} \in \mathsf{U}, \qquad \forall \quad 0 \le s < t < \infty\,, \tag{4}$$

where the rate-set U is precisely the same as given previously for the fluid model.

(iii) The stochastic process $\boldsymbol{A}$ evolves on $\mathbb{R}^\ell$ with overall rate given by

$$\lim_{n\to\infty} \frac{A(tn)}{n} = \alpha t, \qquad a.s.,\ t \ge 0\,.$$

(iv) Each of the stochastic process $\{\boldsymbol{R}, \boldsymbol{S}\}$ are *random functions* from $\mathbb{R}_+^{\ell_u} \to \mathbb{R}_+^\ell$. For any $\zeta \in \mathsf{U}$ the pair of stochastic processes $\{R(\zeta t), S(\zeta t) : t \in \mathbb{R}_+\}$ obey the long-run rate condition,

$$\lim_{r\to\infty} \frac{R(r\zeta t) - S(r\zeta t)}{r} = B\zeta t, \qquad a.s.,\ t \ge 0\,.$$

We note that for each $z \in \mathbb{R}_+^{\ell_u}$, the random variables $S(z)$ and $R(z)$ are typically highly correlated.

The rate condition in (iv) provides a justification of the fluid model (2) through scaling the system equations (1). For $r \ge 1$ define

$$q^r(t;x) = \frac{Q(rt;rx)}{r}, \qquad x \in \mathsf{X},\ t \ge 0\,. \tag{5}$$

Suppose that the open-loop, constant control is applied, $z(t) = \zeta t$, $t \ge 0$, where $\zeta \in \mathsf{U}$ is given. We then have the approximation [6], for any initial $x \in \mathsf{X}$, and any time $t \in \mathbb{R}_+$,

$$\begin{aligned} q^r(t;x) &= x + r^{-1}[-S(Z(rt;rx)) + R(Z(rt;rx)) + A(rt)] \\ &\to x + B\zeta t + \alpha t, \qquad r \to \infty \ a.s. \end{aligned}$$

2.2 Workload relaxations

To define precisely what is meant by *heavy-traffic* we now give a general formulation of network load. It is defined with respect to the fluid model (2) since it is a property of the 'mean flow' of the model. The definition is independent of network variability.

Stabilizability and network load The network model (2) is *stabilizable* if, for any initial condition $x \in \mathsf{X}$, one can find an allocation $\boldsymbol{z}$ and a time $T > 0$ such that $q(T; x) = x + Bz(T) + \alpha T = \boldsymbol{\theta}$. If this is the case, then the network can be controlled so that, starting empty, it will remain empty. Consequently, there exists at least one solution $\zeta^{\rm ss} \in \mathsf{U}$ to the equilibrium equation

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

In the *scheduling problem* the number of activities is equal to the number of buffers in the network. Hence the matrix B is square, and if it is invertible then $\zeta^{\rm ss} = -B^{-1}\alpha$. We then define the *vector load* by

$$\boldsymbol{\rho} = (\rho_1, \ldots, \rho_\ell)^T = C\zeta^{\rm ss} = -CB^{-1}\alpha\,, \tag{6}$$

and the *system load* is defined to be the maximum, $\rho = \max_i \rho_i$. It is clear that $\zeta^{\rm ss} \in \mathsf{U}$ if and only if $\rho \le 1$.

When B is not square then a definition of load requires further effort. First, note that the linear fluid model (2) may be described as a *differential inclusion* on $\mathsf{X} = \mathbb{R}^\ell_+$,

$$\dot{q}(t; x) \in \mathsf{V}, \qquad q(t; x) \in \mathsf{X},$$

where

$$\mathsf{V} := \{v : \langle \xi^i, v\rangle \ge -(\delta_i - \rho_i), \qquad 1 \le i \le \ell_v\}.$$

The vectors $\{(\xi^i, \eta^i)\} \subset \mathbb{R}^\ell \times \mathbb{R}^{\ell_m}_+$ are given in [31] together with several simple examples. The construction given in [31] defines an integer $\ell_r < \ell_v$ such that $\delta_i = 1$ for $i \le \ell_r$, and $\delta_i = 0$ for $i > \ell_r$. For the general network model we define $\rho_i = \langle \xi^i, \alpha\rangle$, $1 \le i \le \ell_v$, and we assume that these load parameters are decreasing in i with $\rho = \rho_1 = \max\{\rho_i\}$.

We let Ξ denote the $\ell_r \times \ell$ matrix with rows equal to $\{\xi^i : 1 \le i \le \ell_r\}$, and for any fluid trajectory $\boldsymbol{q}$ we define the workload process by

$$w(t; w_0) = \Xi q(t; x_0), \qquad t \ge 0\,.$$

Fluid workload relaxation models
For arbitrary $1 \le n \le \ell_r$, the *nth workload relaxation* of (2) is defined as follows.

(i) The state space X is taken as $\mathbb{R}^\ell_+$, and the velocity set is given by

$$\widehat{\mathsf{V}} = \big\{v : \langle \xi^i, v\rangle \ge -(1 - \rho_i), 1 \le i \le n\big\}\,.$$

We denote by $\widehat{\boldsymbol{q}}$ any trajectory in X satisfying $\frac{d^+}{dt}\widehat{q}(t) \in \widehat{\mathsf{V}}$, $t > 0$, where $\frac{d^+}{dt}$ denotes the right-derivative.

(ii) The workload process for the relaxation is given by,

$$\widehat{w}(t; x) = \widehat{\Xi}\widehat{q}(t; x), \qquad t \ge 0,\ x \in \mathsf{X},$$

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

We assume that the workload vectors $\{\xi^i : 1 \le i \le n\}$ are linearly independent. It then follows that the dynamics of $\widehat{w}(t; x)$ are *decoupled*. That is, $\widehat{\boldsymbol{w}}$ is defined as the state process for a differential inclusion with constraints,

$$\frac{d}{dt}\widehat{w}_i(t; x) \ge -(1 - \rho_i), \qquad 1 \le i \le n\,.$$

Equivalently, $\widehat{\boldsymbol{w}}$ is described by the linear system,

$$\widehat{w}_i(t;w) = w - (1-\rho_i)t + \iota_i(t), \qquad 1 \le i \le n\,, \tag{7}$$

where the control $\boldsymbol{\iota} := \{\iota(t) \in \mathbb{R}^n_+ : t \in \mathbb{R}_+\}$ represents *idle-time* of various resources in the network. The control and state process are subject to the following constraints:

$$\begin{aligned} \iota_j(t) - \iota_j(s) &\ \ge\ 0, \qquad t, s \in \mathbb{R}_+,\ t \ge s,\ j = 1, \dots, n. &(8)\\ \widehat{w}(t;x) &\ \in\ \widehat{\mathsf{W}} := \{\widehat{\Xi} x : x \in \mathsf{X}\}. &(9)\end{aligned}$$

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

Probabilistic workload relaxation models A formal workload relaxation can be formulated for the stochastic model through the introduction of an exogenous disturbance $\boldsymbol{N}$:

$$\widehat{W}(t;w) = w - (\boldsymbol{1} - \boldsymbol{\rho})t + I(t) + N(t), \qquad w \in \widehat{\mathsf{W}}\,. \tag{10}$$

The stochastic process $\boldsymbol{N}$ is an n-dimensional Brownian motion with zero-drift and instantaneous covariance $\Sigma > 0$.

The control $\boldsymbol{I}$ again represents cumulative idle-time at various resources in the network. It is constrained to be adapted to the Brownian motion $\boldsymbol{N}$, with the simple constraint $I_i(t) - I_i(s) \ge 0$ for all $t \ge s$, and all $1 \le i \le n$.

The model (10) can be justified through scaling the stochastic model (1) - this is a refinement of the scaling given in (5). However, rather than justify the existence of a limiting RBM model, which is far beyond the scope of this paper, we instead investigate properties of the relaxation (10).

3 Optimal control

Throughout the paper we denote by $c\colon \mathbb{R}^\ell \to \mathbb{R}_+$ a *cost function* on buffer-space. It is assumed that c is a norm on $\mathbb{R}^\ell$. Consequently, it is radially homogeneous, convex, and vanishes only at the origin.

For the fluid model we consider the *total-cost optimal control problem*: For any initial condition $q(0) = x$, we seek an allocation $\boldsymbol{z}$ that minimizes

$$J(x) = \int_0^\infty c(q(t;x))\, dt, \qquad x \in \mathsf{X}\,. \tag{11}$$

We let $J^*(x)$ denote the 'optimal cost', i.e., the infimum over all policies.

For the stochastic model we consider the steady state cost,

$$\gamma = \limsup_{t\to\infty} \mathsf{E}[c(Q(t;x)]\,.$$

Let us consider the problem of optimizing the RBM model. Optimization of the fluid model (2) is considered in detail in [31].

Consider now the model (10) where the exogenous disturbance $\boldsymbol{N}$ has instantaneous covariance $\Sigma > 0$. To effectively translate an optimal control $\widehat{z}_*(t;x)$ for the relaxed fluid model to (1) or (10) one must understand how variability impacts policy structure for the stochastic model.

We assume that there exists a solution (h_*, γ_*) to the following average-cost optimal control equations,

$$h_*(w) = \inf \mathsf{E}\Big[\int_0^T (\bar{c}(\widehat{W}(s;w)) - \gamma_*)\, ds + h_*(\widehat{W}(T;w))\Big], \qquad T \ge 0, \tag{12}$$

where h_* is a continuous function on $\widehat{\mathsf{W}}$, normalized so that $h_*(0) = 0$, and $\gamma_* > 0$ is a constant (see [33, 13]). Typically in average-cost optimization problems the function h_* is also the value-function of the infinite-horizon optimization,

$$h_*(w) = \inf \int_0^\infty \mathsf{E}[\bar{c}(\widehat{W}(t;w)) - \gamma_*]\, dt \tag{13}$$

In either form, the infimum is with respect to all admissible controls $\boldsymbol{I}$. We further assume that, for each $w \in \mathbb{R}^n_+$, there exists an admissible idleness process $\{I^*(t;w) : t \in \mathbb{R}\}$ that achieves the infimum above.

The following result follows as in [29].

Proposition 3.1. *Suppose that a solution (h_*, γ_*) to (12) exists with h_* continuous. Then,*

(i) *γ_* is the optimal average cost: for each $w \in \widehat{\mathsf{W}}$,*

$$\begin{aligned} \gamma_* &= \inf \Big(\lim_{T\to\infty} \mathsf{E}\Big[\tfrac{1}{T}\textstyle\int_0^T \bar{c}(\widehat{W}(t;w))\, dt\Big]\Big) \\ &= \lim_{t\to\infty} \mathsf{E}\big[\bar{c}(\widehat{W}^*(t;w))\big] \end{aligned}$$

where the infimum is over all admissible controls, and the second equation uses the expectation with respect to the optimal control $\boldsymbol{I}^*$.

(ii) *The relative value function h_* is convex and monotone on $\widehat{\mathsf{W}}$.*

(iii) *The relative value function is quadratically bounded:*

$$0 < \liminf_{\|w\|\to\infty} \Big(\frac{h_*(w)}{\|w\|^2}\Big) < \limsup_{\|w\|\to\infty} \Big(\frac{h_*(w)}{\|w\|^2}\Big) < \infty.$$

(iv) *As $n \to \infty$,*

$$\frac{1}{n^2} h_*(nw) \to \widehat{J}_*(w), \qquad w \in \widehat{\mathsf{W}}, w \neq 0.$$

□

Proposition 3.1 tells us that the relative value function is approximated by the fluid value function for large $w \in \widehat{\mathsf{W}}$. One would expect a similar result to hold for models with small variability. We next address this question: *How do optimal control solutions vary with increased system variability?*

Consider the following family of models with increasing variability, parameterized by the real variable $\kappa \geq 0$,

$$\widehat{W}(t;w,\kappa) = w - (\boldsymbol{1} - \boldsymbol{\rho})t + I(t) + \sqrt{\kappa}N(t), \qquad \widehat{W}(0) = w \in \widehat{\mathsf{W}}\,. \tag{14}$$

This is a perturbation of (10). We again assume that $\boldsymbol{N}$ is a drift-less Brownian motion with strictly positive covariance Σ. It follows that the scaled process $N(t;\kappa) := \sqrt{\kappa}N(t)$ is a Brownian motion with instantaneous covariance $\kappa\Sigma$. For $\kappa = 1$ we suppress the dependency of $\widehat{W}$ on κ and write $\widehat{W}(t;w)$ for the process starting from $w \in \widehat{\mathsf{W}}$.

To understand the model (14) we introduce a second parametrized family of network models by scaling the original model (10). For any $\kappa > 0$ define

$$\widehat{W}^\kappa(t;w) = \kappa\widehat{W}(t/\kappa; w/\kappa).$$

This is also described as a linear model

$$\widehat{W}^\kappa(t;w) = w - (\mathbf{1} - \boldsymbol{\rho})t + I^\kappa(t) + N^\kappa(t), \qquad \widehat{W}^\kappa(0) = w \in \widehat{\mathsf{W}}\,, \tag{15}$$

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

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

The following result is immediate from the definitions:

Lemma 3.2. *For any $\kappa > 0$ the following two stochastic processes are identical in law:*

$$N^\kappa(t) = \kappa N(t/\kappa)\,, \qquad N(t;\kappa) = \sqrt{\kappa} N(t) \qquad t \geq 0.$$

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 3.3 also provides a refinement of Proposition 3.1.

We henceforth assume that there is a region $\mathcal{R}^*(\kappa) \subset \widehat{\mathsf{W}}$ that defines the optimal control $\boldsymbol{I}^*$ so that $\widehat{W}(t;w,\kappa) \in \mathcal{R}^*(\kappa)$ a.s. for all $t > 0$ and $\frac{d}{dt}I^*(t) = \boldsymbol{\theta}$ when $\widehat{W}^*(t;w,\kappa)$ lies in the interior of $\mathcal{R}^*(\kappa)$. Assume that a stochastic process $\{\widehat{\boldsymbol{W}}^*, \boldsymbol{I}\}$ satisfying these constraints exists, and that $\widehat{\boldsymbol{W}}^*$ is a strong-Markov process.

Proposition 3.3. *Suppose the assumptions of Proposition 3.1 hold so that a continuous solution (h_*, γ_*) exists to the average cost optimality equations for the model (14) with $\kappa = 1$.*

Then, for any $\kappa > 0$ and any $w \in \widehat{\mathsf{W}}$

(i) $\gamma_*(\kappa) = \kappa\gamma_*(1)$*;*

(ii) $h_*(w;\kappa) := \kappa^2 h_*(w/\kappa; 1)$ *defines the relative value function for (14);*

(iii) *Suppose that $h_*(\,\cdot\,;\,\cdot\,)$ is C^2 in $K \times [0,1]$, where $K \subset \mathcal{R}^*(0)$ is compact. Then, letting $\mathcal{R}_K$ denote the positive cone,*

$$\mathcal{R}_K = \{aw : a > 0, w \in K\}\,,$$

there exists a radially-homogenous function $\ell\colon \mathcal{R}_K \to \mathbb{R}$ such that

$$h_*(w;\kappa) = \widehat{J}_*(w) + \kappa\ell(w) + O(\kappa^2), \qquad w \in \mathcal{R}_K.$$

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

(iv) *For all $\kappa > 0$,*

$$\mathcal{R}^{*'}(\kappa) := \kappa\mathcal{R}^*(1) := \{\kappa w : w \in \mathcal{R}^*(1)\}\,,$$

defines an optimal region for the model (14).

4 Results on Sensitivity in Optimization of RBM Models

This section further develops structural properties of optimal control solutions. Our goal is to obtain bounds on sensitivity of cost with respect to the following variables,

(i) Switching curves defining a policy;

(ii) System variability; and

(iii) Hard constraints on the network.

Throughout most of this section *we restrict to the two dimensional case, and we assume that* $\widehat{\mathsf{W}} = \mathbb{R}^2_+$.

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

4.1 Affine policies

For a two dimensional workload relaxation there are three cases to be considered

Case I The effective cost $\bar{c}$ is monotone on $\mathbb{R}^2_+$.

Case II The effective cost $\bar{c}$ is not monotone on $\mathbb{R}^2_+$, but the vector $(1-\rho_1, 1-\rho_2)^T$ lies in $\widehat{\mathsf{W}}^+$.

Case III The effective cost $\bar{c}$ is not monotone on $\mathbb{R}^2_+$, and $(1-\rho_1, 1-\rho_2)^T \notin \widehat{\mathsf{W}}^+$.

In [31] it is argued that the fluid-model relaxation admits a path-wise optimal solution from any initial condition in Cases I and II: the optimal trajectories evolve in the set $\widehat{\mathsf{W}}^+$ for all $t > 0$

Path-wise optimality cannot hold for any initial conditions $w \notin \widehat{\mathsf{W}}^+$ in Case III. Optimal trajectories evolve in a closed positive cone $\mathcal{R}^*(0)$ that is strictly larger than $\widehat{\mathsf{W}}^+$. It is of the form

$$\mathcal{R}^*(0) = \{w \in \mathbb{R}^2_+ : w_2 \geq s^*_1(w_1; 0), \quad w_1 \geq s^*_2(w_2; 0)\},$$

where the functions s^*_i are linear for each i. The '0' in this notation refers to the assumption that $\kappa = 0$.

When $\kappa > 0$ then we can once again conclude that the optimal solution is trivial, in the sense that $\widehat{\boldsymbol{W}}$ is point-wise minimal, in Case I. This is not true in Cases II or III: the model (10) does not admit a path-wise optimal solution from *any* initial condition.

In Case II we attempt to obtain a solution that is approximately path-wise optimal *in the mean* for $\kappa > 0$. This policy takes on a simple form:

Affine policies

(i) A policy for (14) is called *affine* if the controlled model is a reflected Brownian motion in an affine domain of the form

$$\mathcal{R}(\kappa) = \{w \in \widehat{\mathsf{W}} : \langle \eta^i, w \rangle \geq d_i \kappa\}, \quad 1 \leq i \leq \ell_d,$$

with $\{\eta^i\} \subset \mathbb{R}^2$, $\{d_i\} \subset \mathbb{R}$.

(ii) An *affine translation* of the optimal policy for the fluid model takes the specific form,

$$s_i(w_i;\kappa,\beta) = m_i^*[w_i - \kappa\beta_i]_+, \qquad w_i \ge 0,\ i = 1,2\,, \tag{16}$$

where $0 \le m_1^* < m_2^* \le \infty$ are the slopes of the optimal switching-curves $\{s_i^*(w_i;0) : i = 1,2\}$. Hence, for any $\kappa > 0$, $\beta \in \mathbb{R}_+^2$, this is a simple "affine shift" of the optimal policy for the fluid model.

The motivation for (16) is the following argument: firstly, from Proposition 3.3 (iv) we have $\mathcal{R}^*(\kappa) = \kappa\mathcal{R}^*$ for arbitrary $\kappa > 0$. The fact that $\{s_i^*\}$ define the upper and lower boundaries of $\mathcal{R}^*(\kappa)$ further implies that,

$$s_i^*(w_i;\kappa) = \kappa s_i^*(w_i/\kappa;1), \qquad w_i \in \mathbb{R}_+,\ \kappa > 0\,.$$

Applying Proposition 3.3 again we obtain the following approximations which suggest that the optimal switching curves $\{s_i^*\}$ are asymptotically affine and justify consideration of affine policies:

Proposition 4.1. *Assume that, for some positive* $\{w_i^0 : i = 1,2\}$*, the following derivatives exist and are negative:*

$$d_i^* := \lim_{\kappa\to 0} \frac{s_i^*(w_i;\kappa) - s_i^*(w_i^0;0)}{\kappa w_i^0}\,.$$

Then, the optimal switching curve $s_i^*(w_i;\kappa)$ *is asymptotically affine: For each* $i = 1,2$*,*

$$\begin{aligned} s_i^*(w_i;\kappa) &= s^*(w_i,0) + \kappa d_i^* + O(\kappa^2), \qquad &\kappa \to 0,\ \forall\, w_i > 0;\\ s_i^*(w_i;1) &= s^*(w_i,0) + d_i^* + O(w_i^{-1}), \qquad &w_i \to \infty\,. \end{aligned}$$

□

Having shown that the optimal policy is asymptotically affine, it is natural to attempt to quantify the additional cost incurred when using the best affine policy, rather than the optimal policy.

In Case II this is possible by analysing the one-dimensional *height processes*. For simplicity we present the definition with respect to the lower switching curve. The definition of $\boldsymbol{H}^2$ is completely analogous.

The height process

(i) The *unconstrained* process $X^1(t)$ is the reflected Brownian motion in the domain $\{w \in \mathbb{R}^2 : w_2 \ge m_1^* w_1\}$, with instantaneous covariance $\kappa\Sigma$, drift $-(\boldsymbol{1} - \boldsymbol{\rho})$, and initial condition $X^1(0) = (w_1, m_1^* w_1)^T$.

(ii) The *height process* associated with the lower boundary of $\mathcal{R}(\kappa)$ is the stochastic process defined by $H^1(t) := X_2^1(t) - m_1^* X_1^1(t)$, $t \ge 0$. This is a one-dimensional reflected Brownian motion whose drift and instantaneous covariance are given by,

$$\delta_{H^1} = (\boldsymbol{1} - \boldsymbol{\rho}) \cdot (-m_1^*, 1)^T \qquad \sigma^2_{H^1}(\kappa) = \kappa\sigma^2_{H^1}(1) = \kappa(-m_1^*, 1)\Sigma(-m_1^*, 1)^T$$

In Case II the upper and lower height processes $\{\boldsymbol{H}^i\}$ are each recurrent since $\delta_{H^i} \ge 0$, $i = 1,2$.

In Theorem 4.2,we show that optimization of $\mathsf{E}[\bar{c}(\widehat{W}(t;w,\kappa,\beta))]$ is essentially equivalent to optimizing a cost function on the height process. This reasoning leads to the following parameters,

$$\begin{aligned} \beta_1^* &= \sigma^2_{H^1}(1)(2m_1^*\delta_{H^1})^{-1}\log\Big(1 + \tfrac{|c_2^2|}{|c_2^1|}\Big)\\ \beta_2^* &= \sigma^2_{H^2}(1)(2m_2^*\delta_{H^2})^{-1}\log\Big(1 + \tfrac{|c_1^1|}{|c_1^2|}\Big) \end{aligned} \tag{17}$$

The affine shift $\beta^* \in \mathbb{R}^2_+$ given in (18)

The following theorem shows that it is possible to describe a simple affine policy that is approximately optimal. The optimal affine shifts $\{\beta^*_i\}$ depend on the slopes of the linear functions that make up the piece-wise linear (effective) cost function, and on the load vector $\boldsymbol{\rho}$ in the following way:

$$
\begin{aligned}
\beta^*_1 &= (2m^*_1\delta_{H^1})^{-1}\log\Big(1+\tfrac{|c^2_2|}{|c^1_2|}\Big) \\
\beta^*_2 &= (2m^*_2\delta_{H^2})^{-1}\log\Big(1+\tfrac{|c^1_1|}{|c^2_1|}\Big)
\end{aligned}
\tag{18}
$$

The affine shift $\beta^* \in \mathbb{R}^2_+$ given in (18) is approximately optimal, with exponentially small regret for small κ:

Theorem 4.2. *Suppose that the effective cost $\bar c$ is piece-wise linear, that $\Sigma > 0$, and that $\delta_{H^i} > 0$ for each i. Then, for any value of β, any fixed $0 < T_1 < T_2 < T^*(w)$, and for each $t \in [T_1, T_2]$, $\|w\| \le 1$, if we let $\widehat{W}(t; w, \kappa, \beta)$ denote the resulting workload process initialized at $\widehat{W}(0) = w$ under the affine policy (16), we have,*

$$
\mathsf{E}[\bar c(\widehat{W}(t;w,\kappa,\beta^*))] \le \mathsf{E}[\bar c(\widehat{W}(t;w,\kappa,\beta))] - O(\kappa)\Big(\frac{\|\beta-\beta^*\|^2}{1+\|\beta-\beta^*\|}\Big) + O(\exp(-M/\kappa))\|\beta-\beta^*\| \tag{19}
$$

for some constant $M > 0$ depending on $T_1 > 0$, $T_2 < T^*(w)$. □

There are many possible extensions and interpretations of Theorem 4.2. We give one corollary here. By exploiting the scaling properties of optimal solutions established in Proposition 3.3, and Proposition 4.1, we obtain approximations for $\kappa = 1$:

Theorem 4.3. *Suppose that $\bar c$ is piece-wise linear, that $\Sigma > 0$, and that $\delta_{H^i} > 0$ for each i. Fix $\beta \in \mathbb{R}^2$, $w \in \widehat{\mathsf{W}}$, and $0 < T_1 < T_2 < T^*(w\|w\|^{-1})$. Then,*

(i) *There exists $M > 0$ independent of w such that for all t satisfying $0 < T_1\|w\| \le t \le T_2\|w\| < T^*(w)$,*

$$
\mathsf{E}[\bar c(\widehat{W}(t;w,\kappa,\beta^*))] \le \mathsf{E}[\bar c(\widehat{W}(t;w,\kappa,\beta))] + B_0\|w\|\exp(-M\|w\|)
$$

(ii) *Assume in addition that $\{s^*_i(w_i,\kappa)\}$ are C^1 on $(0,1)\times[0,1]$. There exists $B_0 < \infty$ independent of w such that for each $i = 1, 2$ and each $w_i > 0$,*

$$
|s^*_i(w_i;1) - s_i(w_i;1,\beta^*)| \le \frac{B_0}{w_i}
$$

That is, $m^*_i = d^*_i$ where d^*_i is as given in Proposition 4.1.

□

The results above only apply to models satisfying Case II: In Case I the switching curves are trivial, and independent of κ; in Case III the associated height process is not positive recurrent so the construction of an optimal affine policy is not possible using the approach introduced here. To better understand policies in Case III we now derive sensitivity formulae for cost with respect to control parameters.

4.2 Sensitivity to policy structure

Tp study sensitivity with respect to *policy structure*, suppose that a perturbation of the affine switching curve (16) is given. To understand its effect on the average cost we again consider a parametrized family of policies.

Suppose that $\Delta\colon \mathbb{R}_+ \to \mathbb{R}$ is a bounded, continuous function, and define for any real parameter $0 \le \eta \le 1$,

$$s(w_1; \kappa, \eta) := \kappa s(w_1/\kappa) + \eta\Delta(w_1/\kappa), \qquad w_1 \in \mathbb{R}_+.$$

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

Theorem 4.4 describes the scaling properties, and asymptotic nature, of the average cost when the parameter η is varied.

Theorem 4.4. *Suppose that the switching curve $s(w_1; \kappa, \eta)$ is stabilizing when $\kappa = 1$ and $0 \le \eta \le 1$, and that the steady state cost $\gamma(1, \eta)$ is C^2 on $[0, 1]$. Then,*

(i) *The steady state cost $\gamma(\kappa, \eta)$ satisfies, $\gamma(\kappa, \eta) = \kappa\gamma(1, \eta\kappa^{-1})$, $\kappa \ge 1$, $0 < \eta \le 1$.*

(ii) *The second derivative vanishes as $\kappa \to \infty$,*

$$\frac{\partial^2}{\partial\eta^2}\gamma(\kappa, \eta) = \kappa^{-1}\Big(\frac{\partial^2}{\partial\eta^2}\gamma(1, \eta\kappa^{-1})\Big), \qquad \kappa \ge 1,\ 0 < \eta \le 1$$

□

4.3 Sensitivity to buffer constraints

Theorem 4.4 suggests that there is little reason to devote effort to exactly optimize a workload-relaxation since sensitivity is very low. First-order sensitivity is zero when the switching curve is an interior-point minimizer, and second-order sensitivity vanishes for models with high variability. These theoretical results are plainly illustrated in the numerical results shown in Section 5. This brings us back to the title of this paper - *where does the sensitivity lie?*

We show here that if hard constraints are placed on the network then first order sensitivity with respect to these constraints is non-zero, and bounded from below as system load increases. We consider 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 [11].

Consider first the fluid model. If buffer constraints are imposed then the state space takes the form $\mathsf{X} = \{x : x \ge \boldsymbol{\theta}, x \le \boldsymbol{b}\}$, where $\boldsymbol{b} \in \mathbb{R}^n_+$ (we allow some entries to be infinite). The workload space is equal to $\widehat{\mathsf{W}} = \{\widehat{\Xi} x : x \in \mathsf{X}\}$, and the effective cost $\bar{c}\colon \widehat{\mathsf{W}} \to \mathbb{R}_+$ is the solution to the nonlinear program,

$$\begin{aligned} \bar{c}(w) := \ \mathbf{min} \quad & c(x) \\ \mathbf{subject\ to} \quad & \Xi x = w \\ & x \le \boldsymbol{b} \\ & x \ge \boldsymbol{\theta}. \end{aligned} \tag{20}$$

We again define $\mathcal{X}^*(w)$ to be the optimizing $x \in \mathsf{X}$. If c is piece-wise linear in x then the effective cost is piece-wise linear in the variables $\{w_i, b_j : 1 \le i \le \ell_w, 1 \le j \le \ell\}$. This may be complex for a large network with many constraints, but it can be written explicitly [11].

Computing the sensitivity of cost with respect to a buffer constraint $b_i < \infty$ is straightforward given the formula (20). For the relaxation we have,

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

where $\Gamma_i \geq 0$ is a Langrange-multiplier. The sensitivity for the original fluid-model is identical, after a transient period, since the two trajectories couple after a fixed time, independent of network load [[31, Theorem 15]].

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

$$\frac{\partial}{\partial b_i} \mathsf{E}[\bar{c}(\widehat{W}(t))] = -\Gamma_i \mathsf{P}(\widehat{W}(t) \in \mathsf{W}_{b_i}), \tag{21}$$

where $\mathsf{W}_{b_i} = \{w : \mathcal{X}^*(w)_i = b_i\}$. For example, suppose that the optimal policy for the workload-relaxation is non-idling for one $\kappa \geq 0$. Then the same holds for any $\kappa > 0$, and we may conclude that (21) holds.

We see here that sensitivity with respect to hard constraints is of order one. Consequently, for fixed $\boldsymbol{b}$, the relative cost $\gamma^*(\kappa, \boldsymbol{b}) - \gamma^*(\kappa)$ increases linearly with system variability.

If the policy also changes with b_i then the formula (21) is no longer valid. However, Theorem 4.4 (ii), and the numerical experiments presented in Section 5 all suggest that the sensitivity of cost with respect to the policy in workload space is low. Taking this for granted, the identity (21) suggests several approximations.

5 Numerical results

We conclude with numerical simulations illustrating the results of Section 4 for model shown in Figure 1 with Poisson arrivals and exponential service distributions. We present affine policies in Cases II and III, and we give results from a series of experiments to test sensitivity of cost with respect to control parameters.

An optimal solution for the fluid model is described in workload space by the *linear policy*,

Work resource 2 at maximal rate if $w_2 \geq m_1^* w_1$ and $w_2 \neq 0$

for some constant m_1^*. For a stochastic model we consider affine policies of the specific form,

Work resource 2 at maximal rate if $(w_2 - \bar{w}_2) \geq m_1^*(w_1 - \bar{w}_1)$ and $w_2 \leq \bar{w}_2$ (22)

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 $\rho_2 = \frac{9}{11}$.

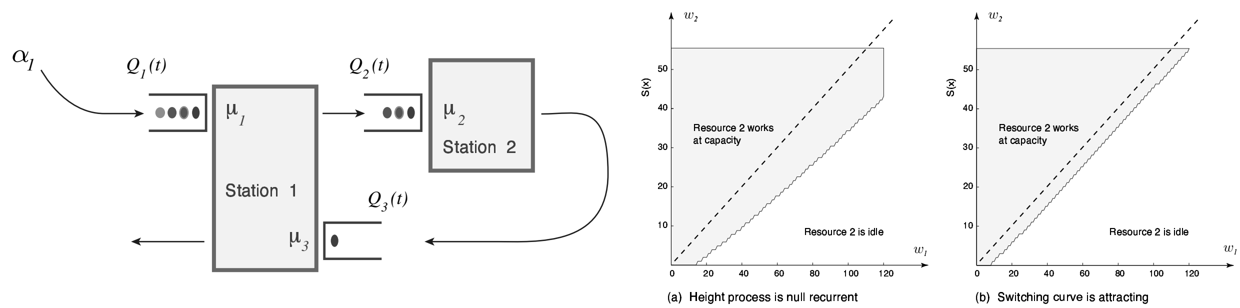

Figure 1: The network used in this section. Optimal policies for the Poisson workload-relaxation in Case II obtained using value iteration.

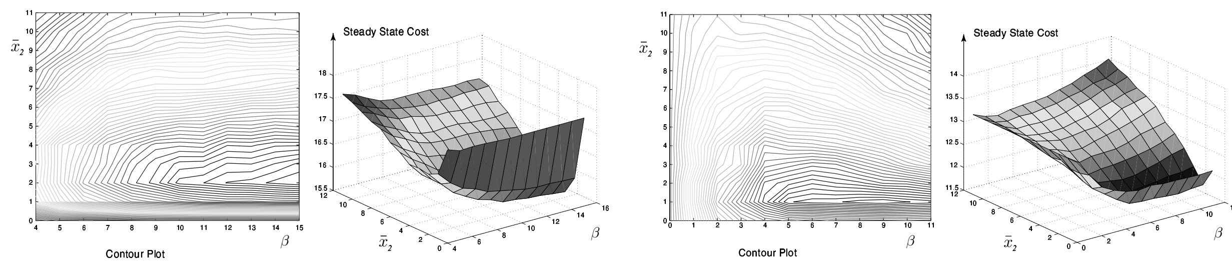

Figure 2: Contour and surface plots of average cost are shown on the left for Case II (a), and on the right for Case II (b), for a sequence of affine translations of the optimal policy for the fluid model.

In each instance the vector $(1-\rho_1, 1-\rho_2)^T$ lies *within* the monotone region $\widehat{\mathsf{W}}^+$. Hence the optimal trajectory for the workload-relaxation $\widehat{\boldsymbol{y}}$ will be *greedy* and *path-wise optimal*. Optimal policies for $\widehat{\boldsymbol{Y}}$ are shown in Figure 1. On the left is Case II (a), where the network is balanced, and on the right is Case II (b). In the balanced case II (a) we see that the difference between the fluid and stochastic switching curves is significant. This corresponds to null-recurrence of the associated *height process*. The optimal policy in Case II (b) is accurately approximated by an affine policy whose offset is determined by an associated Brownian workload model.

We now compare a family of affine policies of the form (22) for the three-dimensional Poisson model. In Case II these may be expressed,

$$\textit{Serve buffer one if} \quad \begin{array}{l} \text{buffer three is zero, or} \\ x_3 \leq \beta \text{ and } x_2 \leq \bar{x}_2. \end{array} \tag{23}$$

For the fluid model, the optimal parameters are $\bar{x}_2 = 0$ and $\beta = \infty$. The simulation results are shown in Figure 2 for Case II (a) and Case II (b).

Numerical results for Case III Now the loads at the two machines are $\rho_1 = \frac{9}{11}$ and $\rho_2 = \frac{9}{10}$ respectively. We are thus in Case III, and 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.

Shown in Figure 3 are the results of a simulation of affine translations of the optimal policies with $\bar{x}_2 = 1, 2, \ldots, 12$ and $\beta = 1, 2, \ldots, 36$, where $\bar{x}_2$ indicates when the control will change and β indicates an affine shift of the switching curve. 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

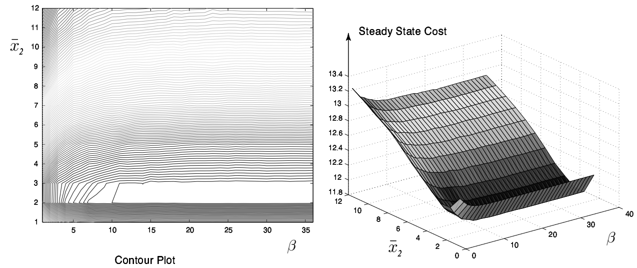

Figure 3: Contour and surface plots of average cost for Case III for a sequence of affine translations of the optimal policy for the fluid model.

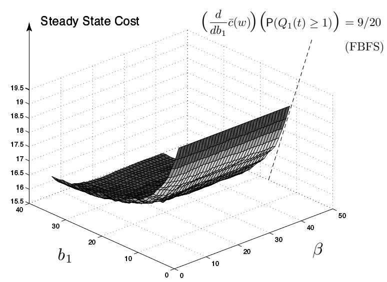

Figure 4: Surface plots of average cost for the balanced model with buffer constraints. The sensitivity estimate (24) is nearly exact in this example. Sensitivity with respect to β remains very small.

the threshold $x_2 \leq \bar{x}_2$ is strictly positive. On the other hand, the sensitivity of the optimal policy to the parameter β is zero for the fluid model at its optimal value $\beta = \bar{x}_3 - \bar{x}_1 - 2\bar{x}_2 = 0$. The contour plot given in Figure 3 shows that this dichotomy is inherited by the stochastic model in this example.

Sensitivity to hard constraints To illustrate the conclusions of Section 4 we focus on buffer constraints. We consider a single numerical experiment in Case II (a). We take $b_2 = b_3 = \infty$, and vary the constraint b_1 on buffer one. We impose the constraint that buffer one receives priority whenever $Q_1(t) \geq b_1$. A policy of this form will maintain a constraint of the form $Q_1(t) \leq b_1 + K$ with high probability (of order $(9/20)^K$). We stress that this is not a loss-model - we do not reject any arriving customers.

A candidate sensitivity approximation with respect to b_1 is

$$\Big[\mathsf{E}^{b_1=k}[c(Q(t))] - \mathsf{E}^{b_1=k-1}[c(Q(t))]\Big] \approx -[\mathsf{P}^{b_1=k}(Q_1(t) \geq k)], \qquad k \geq 1$$

The simplest case is $b_1 = k = 0$ since the resulting policy is First-Come 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 $\mathsf{P}(Q_1(1) \geq 1) = \alpha/\mu_1 = 9/20$, and this gives the approximation,

$$\mathsf{E}^{b_1=1}[c(Q(t))] - \mathsf{E}^{b_1=0}[c(Q(t))] \approx -9/20 \tag{24}$$

Numerical results are shown in the surface plot shown in Figure 4. We see in this example that the approximation (24) is nearly exact.

6 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 parametrized 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 found that increasing variability typically results in increasingly conservative optimal solutions (see Proposition 3.3). 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 [31, 11]). 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 [14].

However, the conclusions of Theorem 4.4 and the numerical results shown in Section 5 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 more detail control issues surrounding buffer constraints; safety-stocks; and non-standard performance metrics such as disaster-recovery [11]. 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 analysed through a sensitivity analysis as described briefly in Section 4.3.

We are also currently investigating a question posed in [31]: *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

References

[1] F. Avram, D. Bertsimas, and M. Ricard. Fluid models of sequencing problems in open queueing networks: an optimal control approach. Technical Report, Massachussetts Institute of Technology, 1995.

[2] 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, New York, 1992.

[3] S.L. Bell and R.J. Williams. Dynamic scheduling of a system with two parallel servers: Asymptotic optimality of a continuous review theshold policy in heavy traffic. In *Proceedings of the 38th Conference on Decision and Control*, pages 1743–1748, Pheonix, Arizona, 1999.

[4] H. Chen and A. Mandelbaum. Discrete flow networks: bottleneck analysis and fluid approximations. *Math. Oper. Res.*, 16(2):408–446, 1991.

[5] H. Chen and D. D. Yao. Dynamic scheduling of a multiclass fluid network. *Operations Research*, 41(6):1104–1115, November-December 1993.

[6] H. Chen and David D. Yao. *Fundamentals of queueing networks: Performance, asymptotics, and optimization*. Springer-Verlag, New York, 2001. Stochastic Modelling and Applied Probability.

[7] R.L. Cruz. A calculus for network delay, part I: Network elements in isolation. *IEEE Trans. Inform. Theory*, 31:114–131, 1991.

[8] J. D. Eng, Humphrey and S. P. Meyn. Fluid network models: Linear programs for control and performance bounds. In J. Cruz J. Gertler and M. Peshkin, editors, *Proceedings of the 13th IFAC World Congress*, volume B, pages 19–24, San Francisco, California, 1996.

[9] J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.*, 5(1):49–77, 1995.

[10] J. G. Dai and S. P. Meyn. Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Trans. Automat. Control*, 40:1889–1904, November 1995.

[11] R. Dubrawski, M. Chen, S. Henderson, and S. P. Meyn. Myopic and far-sighted strategies for control of demand-driven networks. Performance and Stability of Multiclass Queueing Networks, invited session at the INFORMS Annual Meeting, November, 2001.

[12] P. Dupuis and H. Kushner. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, New York, 2001. Applications of Mathematics, Vol 24.

[13] E. A. Feinberg and A. Shwartz. *Handbook of Markov Decision Processes*. Kluwer Academic Publishers, Boston, 2001.

[14] S. B. Gershwin. *Manufacturing Systems Engineering*. Prentice–Hall, Englewood Cliffs, NJ, 1993.

[15] J. M. Harrison. *Brownian Motion and Stochastic Flow Systems*. Wiley, New York, NY, 1985.

[16] J. M. Harrison. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic differential systems, stochastic control theory and applications (Minneapolis, Minn., 1986)*, pages 147–186. Springer, New York, 1988.

[17] J. M. Harrison. Brownian models of open processing networks: Canonical representations of workload. *Ann. Appl. Probab.*, 10:75–103, 2000.

[18] J. M. Harrison. Stochastic networks and activity analysis. In Y. Suhov, editor, *Analytic methods in applied probability. In memory of Fridrih Karpelevich*. American Mathematical Society, Providence, RI, 2002.

[19] J. M. Harrison and L. M. Wein. Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Operations Research*, 38(6):1052–1064, 1990.

[20] F.C. Kelly and C.N. Laws. Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *QUESTA*, 13:47–86, 1993.

[21] S. Kumar and M. Muthuraman. A numerical method for solving singular Brownian control problems. In *Proceedings of the 39th Conference on Decision and Control*, 2000.

[22] H. J. Kushner. *Heavy traffic analysis of controlled queueing and communication networks*. Springer-Verlag, New York, 2001. Stochastic Modelling and Applied Probability.

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

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

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

[26] C. Maglaras. Discrete-review policies for scheduling stochastic networks: Trajectory tracking and fluid-scale asymptotic optimality. *Ann. Appl. Probab.*, 10:–, 2000.

[27] L. F. Martins and H. J. Kushner. Routing and singular control for queueing networks in heavy traffic. *SIAM Journal on Control and Optimization*, 28:1209–1233, 1990.

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

[29] S. P. Meyn. The policy iteration algorithm for average reward Markov decision processes with general state space. *IEEE Trans. Automat. Control*, 42(12):1663–1680, 1997.

[30] S. P. Meyn. Stability and optimization of queueing networks and their fluid models. In *Mathematics of stochastic manufacturing systems (Williamsburg, VA, 1996)*, pages 175–199. Amer. Math. Soc., Providence, RI, 1997.

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

[32] S. P. Meyn. Sequencing and routing in multiclass queueing networks. Part II: Workload relaxations. *SIAM J. Control Optim.*, 2002. To appear.

[33] S. P. Meyn. Multiplicative ergodicity, spectral theory, and phase transitions for Markovian models. Invited presentation in Extracting Low Dimensional Effective Dynamics: Algorithms, Analysis and Applications, Berlin, September, 2001. First SIAM-EMS Conference on Applied Mathematics in Our Changing World, September 2-6, 2001.

[34] J. Perkins. *Control of Push and Pull Manufacturing Systems*. PhD thesis, University of Illinois, Urbana, IL, September 1993. Technical report no. UILU-ENG-93-2237 (DC-155).

[35] M. I. Reiman. Open queueing networks in heavy traffic. *Mathematics of Operations Research*, 9:441–458, 1984.

[36] M. I. Reiman. A multiclass queue in heavy traffic. *Advances in Applied Probability*, 20:179–207, 1988.

[37] A. Shwartz and A. Weiss. *Large deviations for performance analysis: queues, communication and computing*. Chapman-Hall, London, UK, 1995.

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

[39] G. Weiss. A simplex based algorithm to solve separated continuous linear programs. Technical Report,Department of Statistics The University of Haifa Mount Carmel 31905, Israel, 2001.