Workload Models for Stochastic Networks: Value Functions and Performance Evaluation

Sean Meyn *Fellow, IEEE*

***Abstract*— This paper concerns control and performance evaluation for stochastic network models. Structural properties of value functions are developed for controlled Brownian motion (CBM) and deterministic (fluid) workload-models, leading to the following conclusions: Outside of a null-set of network parameters,**

(i) **The fluid value-function is a smooth function of the initial state. Under further minor conditions, the fluid value-function satisfies the derivative boundary conditions that are required to ensure it is in the domain of the extended generator for the CBM model. Exponential ergodicity of the CBM model is demonstrated as one consequence.**
(ii) **The fluid value-function provides a *shadow function* for use in simulation variance reduction for the stochastic model. The resulting simulator satisfies an exact large deviation principle, while a standard simulation algorithm does not satisfy any such bound.**
(iii) **The fluid value-function provides upper and lower bounds on performance for the CBM model. This follows from an extension of recent linear programming approaches to performance evaluation.**

I. Introduction

Several new approaches to performance evaluation and optimization for stochastic network models have been proposed over the past decade. Although the motivation and goals are diverse, a common thread in each of these techniques is the application of an approximate value function, or Lyapunov function:

(i) In the series of papers [1], [2], [3], [4], [5], linear programs are constructed to find a Lyapunov function that will provide upper and lower bounds on steady-state performance. These Lyapunov functions are quadratic, or piecewise-quadratic, since in each of these papers the cost criterion is a linear function of the queue-length process.
(ii) In [6], [7] a quadratic or piecewise-quadratic Lyapunov function is proposed as an initialization in value-iteration or policy-iteration for policy synthesis. An example of a Lyapunov function is the associated *fluid value-function* (see equation (42)). The numerical results reported in [7] show significant improvements when the algorithm is initialized with an arbitrary quadratic Lyapunov function. When initialized with the fluid value-function, convergence is nearly instantaneous. Related ideas are explored in [8].
(iii) The variance-reduction techniques developed in [9], [10], [11] are similarly based on the application of a Lyapunov function. In [10] it is found that the simulator based on the fluid value-function results in dramatic variance reductions for the models considered. Results from one such experiment are shown in Figure 5.

Much of this research has been numerical and algorithmic in nature. In particular, there has been little theory to explain the dramatic numerical results reported in [7], [10]. The purpose of the present paper is to provide a theoretical foundation for analysis of these algorithms. These results are based primarily on new bounds on the dynamic programming equations arising in control formulations for fluid and stochastic network models.

The main results of this paper concern workload models for the network of interest. This viewpoint is motivated by recent results establishing strong solidarity between workload models and more complex queueing model in buffer-coordinates [12], [13], [14], [15], [16], [17], [18], [19]. In particular, under appropriate statistical assumptions and appropriate assumptions on the operating policy, a Gaussian or deterministic workload model is obtained as a limit upon scaling the equations for a standard stochastic queueing model (see [20], [21], [22], [18]).

The starting point of this paper is similar to the approach used in [23], [24] to establish positive recurrence for a reflected diffusion. The authors first construct a Lyapunov function for a deterministic fluid model, and then show that the same function or a smoothed version serves as a Lyapunov function for the diffusion (i.e., condition (V2) of [25], [26] is satisfied.)

The present paper considers both controlled fluid models and controlled Brownian motion (CBM) models for the workload process. The main conclusions are,

(i) Structural properties of the fluid value-function J are developed: under mild conditions, the value-function is C^1, and its directional derivatives along the boundaries of the 'constraint region' are zero, when the direction is taken as the particular reflection vector at this point on the boundary (see Theorem 4.2.)
(ii) These structural properties imply that the fluid value-function is in the domain of the extended generator of the CBM workload model. As one corollary, it is shown in Theorem 4.7 that the stochastic workload process satisfies a strong form of geometric ergodicity for the policies considered.
(iii) It is shown in Theorem 5.3 that the smoothed estimator of [10] based on the function J satisfies an exact large deviation principle. It is argued that the standard simulation algorithm does not satisfy any such bound.

Coordinated Science Laboratory and the Department of Electrical and Computer Engineering at the University of Illinois, 1308 W. Main Street, Urbana, IL 61801. URL decision.csl.uiuc.edu:80/$\sim$meyn (meyn@uiuc.edu). This paper is based upon work supported by the National Science Foundation under Award Nos. ECS 02 17836 and DMI 00 85165.

(iv) The linear-programming approaches of [1], [2], [3], [4], [5] are extended based on the structural results obtained in Theorem 4.2.

To provide motivation, we begin in Section II with a description of various models for a simple 2-station network. A short survey of control and performance evaluation techniques is presented, focusing on this single example.

General workload models are described in Section III, along with basic properties of the associated value functions. These structural results are refined and applied in Section IV to provide stability results for the model, Section V contains applications to simulation, and examples illustrating linear-programming approaches are described in Section VI. Conclusions and topics of future research are included in Section VII.

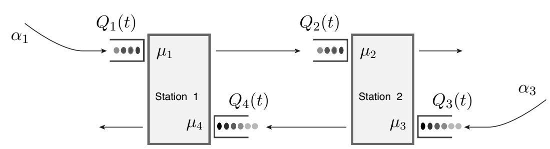

Fig. 1. The Kumar-Seidman-Rybko-Stolyar (KSRS) network.

II. An Illustrative Example

The *Kumar-Seidman-Rybko-Stolyar* (KSRS) network shown in Figure 1 will be used to develop terminology and illustrate the issues to be considered in the remainder of the paper. Further details regarding fluid models, random walk models, Brownian models, and the relationships between them may be found in [12], [27], [14], [16], [10], [28], [17], [18], [29], [8].

A. Models for the queue-length process

The network shown in Figure 1 has two stations, and four buffers. A convenient model is described in discrete-time as follows. We let $\boldsymbol{Q}$ denote the queue-length process, and $\boldsymbol{U}$ denote the allocation process, both evolving on $\mathbb{Z}_+^4$.[1] The allocation process is subject to linear constraints,

$$U(k) \in \mathsf{U} := \{\zeta \in \mathbb{R}_+^4 : C\zeta \le 1,\ \zeta \ge 0\}, \tag{1}$$

where $0 \in \mathbb{R}^4$ denotes a vector of zeros, and inequalities between vectors are interpreted component-wise. The 2×4 *constituency matrix* is given by

$$C := \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{pmatrix}.$$

Let $(\boldsymbol{A}, \boldsymbol{S})$ denote the service and arrival processes at the various buffers. For each $i = 1, \dots, 4$, and each time $k \ge 1$, the distribution of $S_i(k)$ is assumed to be Bernoulli, and the distribution of $A_i(k)$ is supported on $\mathbb{Z}_+$. We assume moreover that $\mathsf{E}[\|A(k)\|^2] < \infty$. The topology shown in Figure 1 implies that $\boldsymbol{A}$ is of the specific form,

$$A(k) = (A_1(k),\ 0,\ A_3(k),\ 0)^{\tau}, \qquad k \ge 1.$$

[1]Throughout the paper we use boldface to denote a vector-valued function of time, either deterministic or stochastic.

The common mean of the arrival and service variables are expressed in the usual notation, $\mathsf{E}[A_i(k)] = \alpha_i$, $\mathsf{E}[S_i(k)] = \mu_i$, $1 \le i \le 4$,

The routing matrix for this model is defined by $R_{ij} = 1$ if customers departing buffer i immediately enter buffer j. Hence in this example, $R_{12} = R_{34} = 1$ and $R_{ij} = 0$ otherwise.

For a given initial queue-length vector $Q(0) \in \mathbb{Z}_+^4$, the queue length at future times is defined recursively via,

$$Q(k+1) = Q(k) + B(k+1)U(k) + A(k+1), \qquad k \ge 0, \tag{2}$$

where $B(k) := -[I - R^{\tau}]\mathrm{diag}(S(k))$, $k \ge 1$. This is a version of the *controlled random walk* (CRW) model considered in [10], [17].

The *fluid model* associated with (2) is the linear system,

$$q(t) = x + Bz(t) + \alpha t, \qquad t \ge 0,$$

where $B := \mathsf{E}[B(k)]$, $\alpha := \mathsf{E}[A(k)]$, $q(0) = x \in \mathbb{R}_+^4$ is the initial condition, and $\boldsymbol{z}$ is the cumulative allocation process evolving on $\mathbb{R}_+^4$. The fluid allocation process is subject to the linear constraints,

$$z(0) = 0, \qquad \frac{z(t) - z(s)}{t - s} \in \mathsf{U}, \qquad t > s \ge 0. \tag{3}$$

In the fluid model the lattice constraints on the queue-length process and allocation process are relaxed. Each is deterministic and Lipschitz continuous as a function of time.

These processes are differentiable a.e. since they are Lipschitz continuous. Letting $\zeta(t) := \frac{d^+}{dt} z(t)$ denote the right derivative of the allocation process, the fluid model can be expressed as the state-space model,

$$\frac{d^+}{dt} q(t) = B\zeta(t) + \alpha, \qquad t \ge 0. \tag{4}$$

The allocation-rate vector $\zeta(t)$ is constrained as in (1), with $\zeta(t) \in \mathsf{U}$ for all $t \ge 0$.

B. Workload models

Throughout much of this paper we restrict attention to relaxations of the fluid model based on workload, and analogous workload models to approximate a CRW stochastic network model.

The two workload vectors associated with the two stations in the KSRS model are defined as follows, where $m_i := \mu_i^{-1}$ for $i = 1, 2, 3, 4$:

$$\xi^1 = (m_1,\ 0,\ m_4,\ m_4)^{\tau}, \quad \xi^2 = (m_2,\ m_2,\ m_3,\ 0)^{\tau}.$$

The two load parameters are given by $\rho_i := \langle \xi^i, \alpha \rangle$, $i = 1, 2$, or $\rho_1 = m_1\alpha_1 + m_4\alpha_3$, $\rho_2 = m_2\alpha_1 + m_3\alpha_3$, and we let $\rho = (\rho_1, \rho_2)^{\tau} \in \mathbb{R}_+^2$.

The workload vectors are used to construct the minimal draining time T^* for the network, thereby solving the time-optimal control problem. Assuming that the load condition holds, $\rho_\bullet := \max(\rho_1, \rho_2) < 1$, it is known that $T^*(x) < \infty$ for each initial condition $x \in \mathbb{R}_+^4$, and we have the following explicit representation,

$$T^*(x) = \max_{1=1,2} \frac{\langle \xi^i, x \rangle}{1 - \rho_i}. \tag{5}$$

A cumulative allocation process $\boldsymbol{z}^*$ is called time-optimal if the resulting state trajectory $\boldsymbol{q}^*$ satisfies $q^*(T) = 0$ in minimum time $T = T^*$. One solution to the time-optimal control problem when $q(0) = x$ is given by $z^*(t) = \zeta^* t$, $0 \le t < T^*$, where $\zeta^* = -B^{-1}[x/T^*(x) + \alpha]$. The optimal state trajectory $\boldsymbol{q}^*$ travels towards the origin in a straight line. This policy and related policies for a stochastic model are examined in [30].

The central motivation for consideration of the workload vectors is that they provide the following definition of the workload process. We let Ξ denote the 2×4 matrix whose rows are equal to the workload vectors $\{\xi^i : i = 1, 2\}$. Equivalently, $\Xi = C\,\mathrm{diag}(\mu_i^{-1})[I - R^{\scriptscriptstyle T}]^{-1} = -CB^{-1}$. Given any state trajectory $\boldsymbol{q}$, the workload process is defined by $w(t) = \Xi q(t)$ for $t \ge 0$.

We have the following simiple representation of $\boldsymbol{w}$ in terms of the two-dimensional *drift vector* defined as $\delta := 1 - \rho$, and the two-dimensional cumulative *idleness process*, defined by $I(t) = t1 - Cz(t)$, $t \ge 0$. For a given initial condition $w = \Xi x \in \mathbb{R}^2_+$, and any $t \ge 0$,

$$w(t) = w + \Xi(Bz(t) + \alpha t) = w - Cz(t) + \rho t = w - \delta t + I(t).$$

The idleness process is non-negative and has non-decreasing components since $\boldsymbol{z}$ satisfies (3).

A relaxation of the fluid model is obtained on relaxing constraints on the idleness process. Consider the two-dimensional model $\widehat{\boldsymbol{w}}$ satisfying $\widehat{w}(0) = w$ and,

$$\widehat{w}(t) = w - \delta t + I(t), \quad t \ge 0. \tag{6}$$

The idleness process $\boldsymbol{I}$ is assumed non-negative, with non-decreasing components, but its increments may be unbounded. We say that $\widehat{\boldsymbol{w}}$ is *admissible* if it is a measurable function of time satisfying (6) with $\boldsymbol{I}$ non-decreasing, $I(0) = 0$, and $\widehat{\boldsymbol{w}}$ restricted to $\mathbb{R}^2_+$.

Relaxations of the stochastic CRW model are developed in [17], [19], following the one-dimensional relaxation introduced in [31]. Here we consider the CBM model based on $\widehat{\boldsymbol{w}}$ through the introduction of an additive disturbance: The two-dimensional process $\widehat{\boldsymbol{W}}$ with initial condition $\widehat{W}(0) = w$ satisfies the equations,

$$\widehat{W}(t) = w - \delta t + I(t) + N(t), \quad \widehat{W}(0) = w \in \mathbb{R}^2_+,\ t \ge 0, \tag{7}$$

where $\boldsymbol{I}$ is again assumed non-negative, with non-decreasing components, and $\widehat{\boldsymbol{W}}$ is restricted to $\mathbb{R}^2_+$. The process $\boldsymbol{N}$ is a 2-dimensional Brownian motion with covariance $\Sigma > 0$.

C. Control

In constructing policies to determine the idleness processes for either the fluid or CBM workload model, we restrict to the following *affine policies* introduced in [17]. For a given convex, polyhedral region $\mathcal{R} \subset \mathbb{R}^2_+$, the corresponding affine policy is defined for the fluid model as follows: For each initial condition $w \in \mathbb{R}^2_+$ we take $\widehat{w}(0) = w$, and for $t > 0$,

$$\widehat{w}(t) = \{\text{minimal element } w' \in \mathcal{R} \text{ satisfying } w' \ge w - \delta t\}.$$

The trajectory $\widehat{\boldsymbol{w}}$ is continuous and piecewise linear on $(0, \infty)$, and exhibits a jump at time $t = 0$ if $w \in \mathcal{R}^c$. The *work-conserving* policy is simply the affine policy with $\mathcal{R} = \mathbb{R}^2_+$. The main result of [17] shows that an optimal policy for the CBM model is approximated by an affine policy when the state is large.

Consider for example the convex polyhedron,

$$\mathcal{R} = \{w \in \mathbb{R}^2_+ : w_2 \ge \theta w_1 - \beta,\ w_1 \ge \theta w_2 - \beta\}, \tag{8}$$

where $\beta \in \mathbb{R}_+$, and $\theta \in [0, 0.5)$ are given. Suppose that $\delta_1 = \delta_2 > 0$, and consider the initial condition $w = (w_1, 0)^{\scriptscriptstyle T}$ for some $w_1 > \beta$. Then $\widehat{w}(0+) = (w_1, \theta w_1 - \beta)^{\scriptscriptstyle T}$, and subsequently, the trajectory $\widehat{\boldsymbol{w}}$ is piecewise-linear on $\mathcal{R}$.

The definition of the affine policy is identical for the CBM model, except of course the trajectories $(\widehat{\boldsymbol{W}}, \boldsymbol{I})$ are not piecewise linear. A general construction is provided in Section III-B.

Let $c\colon \mathbb{R}^4 \to \mathbb{R}_+$ denote the ℓ_1-norm on $\mathbb{R}^4$. A typical objective function in the queueing theory literature is the steady-state cost for the CRW model, defined as,

$$\eta = \limsup_{k\to\infty} \mathsf{E}_x[c(Q(k))], \qquad x \in \mathbb{Z}^4_+. \tag{9}$$

This is independent of x for 'reasonable' policies, including the average-cost optimal policy [6], [32]. The optimal policy is stationary and Markov (the action $U(k)$ is a function of the current state $Q(k)$), so that $\boldsymbol{Q}$ is a time-homogeneous Markov chain [6].

Given this cost function on buffer-levels, the *effective cost* for the two-dimensional relaxation is defined as the value of the linear program,

$$\begin{array}{rlrcl} \bar{c}(w) = & \textbf{min} & x_1 + x_2 + x_3 + x_4 & & \\ & \textbf{s. t.} & \langle \xi^1, x\rangle & = & w_1 \\ & & \langle \xi^2, x\rangle & = & w_2,\ x \in \mathbb{R}^4_+. \end{array} \tag{10}$$

That is, $\bar{c}(w)$ is the cost associated with the 'cheapest' state $x \in \mathbb{R}^4_+$ satisfying the given workload values. Further motivation is provided in [13], [17], [16], [10], [28]. The effective cost is piecewise linear: Letting $\{\bar{c}^i\}$ denote the extreme points in the dual of the linear program (10), we have $\bar{c}(w) = \max_i \langle \bar{c}^i, w\rangle$ for $w \in \mathbb{R}^2_+$.

For a given affine policy the associated value function is defined by,

$$J(w) = \int_0^\infty \bar{c}(\widehat{w}(t))\,dt, \qquad \widehat{w}(0) = w \in \mathbb{R}^2_+.$$

It follows from the definition that the dynamic programming equation holds,

$$\int_0^T \bar{c}(\widehat{w}(t))\,dt + J(\widehat{w}(T)) = J(\widehat{w}(0)), \qquad T \ge 0. \tag{11}$$

Suppose that J is differentiable at some w within the interior of $\mathcal{R}$. Then on setting $\widehat{w}(0) = w$, dividing each side of (11) by T, and letting $T \downarrow 0$ we obtain the differential dynamic programming equation,

$$\bar{c}(w) + \langle \nabla J(w), -\delta\rangle = 0. \tag{12}$$

The fluid value function is constructed in Proposition 2.1 based on (12) in the following two special cases:

CASE I $\quad \mu_2 = \mu_4 = 1/3$ and $\mu_1 = \mu_3 = 1$:

$$\bigl\{\bar{c}^i : i = 1, 2, 3\bigr\} = \bigl\{\tfrac{1}{3}\tbinom{1}{0},\ \tfrac{1}{4}\tbinom{1}{1},\ \tfrac{1}{3}\tbinom{0}{1}\bigr\}$$

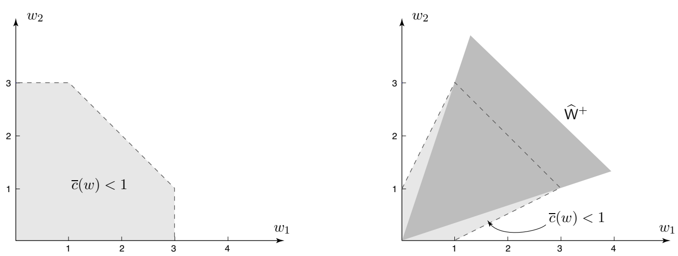

Fig. 2. Level sets of the effective cost for the KSRS model in Cases I and II respectively under the work-conserving policy.

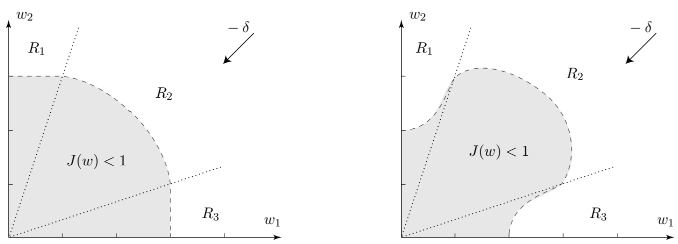

Fig. 3. Level sets of the value function J for the KSRS model.

CASE II $\mu_2 = \mu_4 = 1$ and $\mu_1 = \mu_3 = 1/3$:

$$\{\bar c^i : i = 1,2,3\} = \{\binom{1}{-2},\ \tfrac{1}{4}\binom{1}{1},\ \binom{-2}{1}\}$$

In each case, the effective cost $\bar c\colon \mathbb{R}^2_+ \to \mathbb{R}_+$ is continuous, piecewise linear, and strictly linear on each of the three regions $\{R_i : i = 1,2,3\}$ shown in Figure 3. Level sets of $\bar c$ and J are shown in Figures 2 and 3, respectively. The level sets shown in Figure 3 are smooth since J is C^1:

Proposition 2.1: Suppose that the 2-dimensional workload relaxation on $\mathbb{R}^2_+$ is controlled using the work-conserving policy. Then, in each of Cases I and II, the value function $J\colon \mathbb{R}^2_+ \to \mathbb{R}_+$ is C^1 and piecewise-quadratic, and purely quadratic within each of the regions $\{R_i : i = 1,2,3\}$. Hence, $J(w) = \frac{1}{2} w^{\scriptscriptstyle T} D^i w$ for $w \in R_i$, $i = 1,2,3$, where in Case I the 2×2 matrices are given by, $\{D^1, D^2, D^3\} =$

$$\Big\{ \tfrac{1}{3}\delta_1^{-1}\begin{bmatrix}1 & 0\\ 0 & 0\end{bmatrix},\quad \tfrac{1}{8}\delta_1^{-1}\begin{bmatrix}3 & -1\\ -1 & 3\end{bmatrix},\quad \tfrac{1}{3}\delta_1^{-1}\begin{bmatrix}0 & 0\\ 0 & 1\end{bmatrix} \Big\} \tag{13}$$

PROOF: The C^1 property is established for general workload models in Theorem 4.2. Computation of J is performed as follows: The value function satisfies the dynamic programming equation (12), or equivalently,

$$-w^{\scriptscriptstyle T} D^i \delta = -\langle \bar c^i, w\rangle, \quad w \in R_i,\ i = 1,2,3.$$

It follows that $D^i\delta = \bar c^i$ for each i. A single constraint on each of $\{D^1, D^3\}$ is obtained on considering w on the two boundaries of $\mathbb{R}^2_+$. Finally, the C^1 property implies that $D^1\binom{1}{3} = D^2\binom{1}{3}$ and $D^3\binom{3}{1} = D^2\binom{3}{1}$. Combining these linear constraints yields (13). □

In Case I the effective cost is monotone, so that $\bar c^i \geq 0$ for each i. In this case the work-conserving policy is *pathwise optimal* for the fluid model, in the sense that $\bar c(\widehat w(t))$ is minimal for each t. It follows that the value function J is convex in this case, which is seen in the plot at left in Figure 3.

In Case II the monotone region for the effective cost is given by $\mathsf{W}^+ = \text{closure}\,(R_2)$, as shown at right in Figure 2. In this case, a pathwise optimal solution exists, but it is defined by the affine policy (8) using $\beta = 0$ and $\theta = \frac{1}{3}$, so that $\mathcal{R} = \mathsf{W}^+$. In particular, for an initial condition of the form $\widehat w(0) = (w_1, 0)^{\scriptscriptstyle T}$, the optimal trajectory is given by $\widehat w(0+) = w_1(1, \frac{1}{3})^{\scriptscriptstyle T}$, and $\widehat w(t) = (w_1 - \delta_1 t)(1, \frac{1}{3})^{\scriptscriptstyle T}$ for $0 < t \leq w_1/\delta_1$.

Under the optimal policy in Case II, the value function is purely quadratic on W^+:

$$J(w) = \tfrac{1}{2} w^{\scriptscriptstyle T} D w; \qquad D = \tfrac{1}{8}\delta_1^{-1}\begin{bmatrix}3 & -1\\ -1 & 3\end{bmatrix}. \tag{14}$$

We now turn to the CBM model (7). In Case I we find that the work-conserving policy in which $\mathcal{R} = \mathbb{R}^2_+$ is again pathwise optimal. This policy defines a reflected Brownian motion on $\mathbb{R}^2_+$ with reflection normal at each boundary.

In Case II a pathwise optimal solution does not exist. The optimal solution evolves in a region $\mathcal{R}^*$ that is a strict subset of $\mathbb{R}^2_+$. In the discounted-cost case it is shown in [17] that the constraint region is approximately affine, in the sense that for large values of w the region $\mathcal{R}^*$ is approximated by an affine domain $\mathcal{R}$ of the form (8) with $\theta = \frac{1}{3}$. Theorems 4.6 and 4.7 and numerical results contained in [17] strongly suggest that this result carries over to the average-cost optimal control problem.

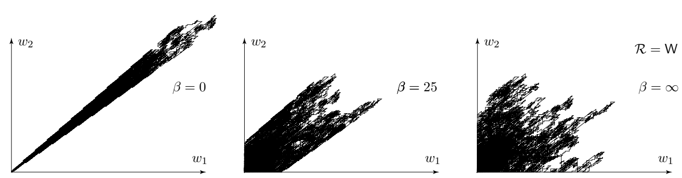

Fig. 4. Simulation of a stochastic workload model using an affine policy. In each case, the constraint region $\mathcal{R}$ was defined by (8) with $\theta = 0.9$. The constant β was defined from left to right by $\beta = 0$, 25, and ∞. The sample paths of $(\boldsymbol{A}, \boldsymbol{S})$ and the initial condition $\widehat W(0) = (150, 0)^{\scriptscriptstyle T}$ were identical in each of the three experiments.

A formula for the parameter β_* in this approximation is also obtained in [17]. Again in Case II, suppose that for some $\sigma^2 > 0$, and $a \in [-1, 1]$ we have $\Sigma_{11} = \Sigma_{22} = \sigma^2$ and $\Sigma_{12} = \Sigma_{21} = a\sigma^2$. Then, Equation (43) of [17] gives,

$$\beta_* = \tfrac{1}{2}\frac{\sigma_H^2}{\delta_H}\ln\Big(1 + \frac{c_+}{c_-}\Big), \tag{15}$$

with $c_+ = \bar c_2^2 = \frac{1}{4}$, $c_- = |\bar c_2^1| = 2$, and

$$\delta_H = (1-\theta)\delta_1, \qquad \sigma_H^2 = (1 - 2a\theta + \theta^2)\sigma^2$$

These are the diffusion parameters for the *height process* defined by $H(t) := W_2(t) - \theta W_1(t)$, $t \geq 0$.

Although (15) is in general an approximation, it defines the average cost optimal policy in some specific cases: Consider the one-dimensional model of [33], and the multidimensional examples treated in [34] where an affine policy with parameters defined using a version of (15) is precisely average-cost optimal.

Figure 4 illustrates the role of the affine shift parameter in a simulation using a CRW approximation of the CBM model. The constraint region used in the figure at left is similar to that defining the optimal policy for the fluid model in an example considered in [28].

A policy constructed for the workload model must finally be translated to form a useful policy for the physical network. Given an affine policy of the form described here using the constraint region $\mathcal{R}$ defined in (8), and a vector $\bar{w} \in \mathbb{R}^2_+$ of *safety-stock* values, a policy for the four-buffer queueing network is described as follows:

Serve $Q_1 \geq 1$ at Station I if and only if $Q_4 = 0$, or

$$\Xi Q \in \mathcal{R} \quad \textit{and} \quad \mu_2^{-1} Q_2 + \mu_3^{-1} Q_3 \leq \bar{w}_2. \qquad (16)$$

An analogous condition holds at Station II.

The safety-stock values are necessary to prevent idleness at each station as $W(k) = \Xi Q(k)$ evolves in $\mathcal{R}$, so that the workload process mimics the CBM model with constraint region $\mathcal{R}$.

D. Performance evaluation

Given a Markov policy for the network such as (16), the controlled process $\boldsymbol{Q}$ is a time-homogeneous Markov chain. We let π denote its unique steady-state distribution, when it exists. In this case we have $\eta = \pi(c) = \sum_{y \in \mathbb{Z}^4_+} \pi(y) c(y)$ for a.e. $x \in \mathbb{Z}^4_+$ $[\pi]$. The steady-state distribution is not easily computed, so we seek bounds or approximations for the steady-state cost η.

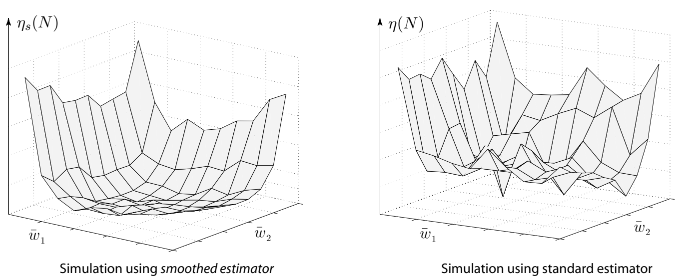

Fig. 5. Estimates of the steady-state customer population in the KSRS model as a function of 100 different safety-stock levels using the policy (16). Two simulation experiments are shown, where in each case the simulation runlength consisted of $N = 200{,}000$ steps. The left hand side shows the results obtained using the smoothed estimator (18); the right hand side shows results with the standard estimator (17).

The most common approach to performance evaluation of a scheduling policy is through simulation. The *standard estimator* of η is defined by the sample-path average,

$$\eta(N) = \frac{1}{N} \sum_{k=0}^{N-1} c(Q(k)), \qquad N \geq 1. \qquad (17)$$

This is consistent in the sense that $\eta(N) \to \eta$ a.s. for a.e. initial condition $x \in \mathbb{Z}^4_+$ $[\pi]$. However, it is well known that long simulation runlengths are required for an accurate estimate of η when the system load is near unity (see [35], [36], and Section V for general results pertaining to the CBM model.)

A version of the *smoothed estimator* of the form considered in [10], [11] is described as follows. For any function G on $\mathbb{R}^4_+$ define $\Delta_G(x) = \mathsf{E}[G(Q(k+1)) - G(Q(k)) \mid Q(k) = x] = \mathsf{E}_x[G(Q(1)) - G(Q(0))]$, $x \in \mathbb{Z}^4_+$. The steady-state mean of Δ_G is zero, provided G is absolutely integrable w.r.t. π. The smoothed estimator is then defined as the sequence of asymptotically unbiased estimates given by,

$$\eta_s(N) = \frac{1}{N} \sum_{k=0}^{N-1} \big(c(Q(k)) + \Delta_G(Q(k)) \big), \quad N \geq 1. \quad (18)$$

The function Δ_G is an example of a *shadow function*, since it is meant to eclipse the function $\bar{c}$ to be simulated [10]. This is a special case of the control variate method [37], [38]. However, the homogeneous structure of the workload model is exploited to obtain a highly specialized estimator based on the fluid value-function and the dynamic programming equation (12) for $\widehat{\boldsymbol{w}}$.

Consider the four dimensional queueing model (2) controlled using the policy (16) under the assumptions of Case I, and let $J \colon \mathbb{R}^2_+ \to \mathbb{R}_+$ denote the fluid value function for $\widehat{\boldsymbol{w}}$ using the same constraint region $\mathcal{R}$. We choose $G \colon \mathbb{R}^4_+ \to \mathbb{R}$ as the C^1, piecewise-quadratic function given by $G(x) = J(\Xi x)$ for $x \in \mathbb{R}^4_+$.

Bounds on the summand in (18) are obtained as follows. By the Mean Value Theorem applied to J we have for some $\vartheta \in (0,1)$, with $W^\vartheta(k) = (1-\vartheta) W(k+1) + \vartheta W(k)$,

$$\begin{aligned} G(Q(k+1)) - G(Q(k)) &= J(W(k+1)) - J(W(k)) \\ &= \langle \nabla J\,(W^\vartheta(k)), W(k+1) - W(k) \rangle \\ &= \langle \nabla J\,(W(k)), W(k+1) - W(k) \rangle \\ &\quad + \langle \nabla J\,(W^\vartheta(k)) - \nabla J\,(W(k)), W(k+1) - W(k) \rangle. \end{aligned}$$

This identity combined with Lipschitz continuity of the function ∇J on $\mathbb{R}^2_+$, and the second-moment assumption on $\boldsymbol{A}$, implies the following bound: Writing $\Delta_W(x) = \mathsf{E}_x[W(1) - W(0)\rangle]$, we have for some $k_0 < \infty$ and all $x \in \mathbb{Z}^4_+$,

$$\Big| \Delta_G(x) - \langle \nabla_w J\,(\Xi x), \Delta_W(x) \rangle \Big| \leq k_0.$$

The policy (16) is designed to enforce $\Delta_W(x) = -\delta$ whenever $w = \Xi x$ lies in $\mathcal{R}$, and when this holds the dynamic programming equation (12) implies that $\langle \nabla_w J\,(\Xi x), \Delta_W(x) \rangle = -\bar{c}(w)$. The policy is also designed to maintain $c(x) \approx \bar{c}(w)$. Consequently $\Delta_G(x) \approx -c(x)$, and this is precisely the motivation for the smoothed estimator.

Shown in Figure 5 are estimates of the steady-state customer population in Case I for the family of policies (16), indexed by the safety-stock level $\bar{w} \in \mathbb{R}^2_+$, with $\mathcal{R} = \mathbb{R}^2_+$. Shown at left are estimates obtained using the smoothed estimator (18) based on this function G, and the plot at right shows estimates obtained using the standard estimator (17). This and other numerical results presented in [10] show that significant improvements in estimation performance are possible with a carefully constructed simulator.

We now develop theory for general workload models to explain these numerical findings.

III. Controlled Brownian Motion Model

In Section III-A a general fluid model is constructed in which the state process is defined as a differential inclusion on $\mathbb{R}^\ell_+$. This is the basis of the general Controlled Brownian Motion workload model defined in Section III-B, and developed in the remainder of the paper.

A. Workload for a fluid model

The state process for the fluid model $\boldsymbol{q}$ has continuous sample paths, and evolves on $\mathbb{R}^\ell_+$. For simplicity we do not treat buffer constraints here [17], [28]. Rate constraints on the sample paths are specified as follows. It is assumed that a bounded, convex, polyhedral *velocity set* $\mathsf{V} \subset \mathbb{R}^\ell$ is given, and that $\boldsymbol{q}$ satisfies,

$$(t-s)^{-1}(q(t)-q(s)) \in \mathsf{V}, \quad 0 \le s < t < \infty. \tag{19}$$

The right derivative $v(t) = \frac{d^+}{dt}q(t)$ exists for a.e. t since $\boldsymbol{q}$ is Lipschitz continuous. This velocity vector is interpreted as a control, subject to the constraint $v(t) \in \mathsf{V}$ for $t \ge 0$.

It is assumed that the fluid model is *stabilizable*. That is, for each initial condition $x^0 \in \mathbb{R}^\ell_+$, one can find $\boldsymbol{v}$ evolving in V and $T < \infty$ such that $q(t) = 0$ for $t \ge T$. We infer that $0 \in \mathsf{V}$ under the stabilizability assumption.

Since V is a convex polyhedron containing the origin, it can be expressed as the intersection of half-spaces as follows. For some integer $n_v \ge 1$, vectors $\{\xi^i : 1 \le i \le n_v\} \subset \mathbb{R}^\ell$, and constants $\{\delta_i : 1 \le i \le n_v\} \subset \mathbb{R}_+$,

$$\mathsf{V} = \{v : \langle \xi^i, v\rangle \ge -\delta_i, \qquad 1 \le i \le n_v\}. \tag{20}$$

These vectors generalize the workload vectors constructed in the KSRS model.

To obtain a relaxation we fix an integer $n < n_v$, and relax the linear constraints in the definition (20) for $i > n$. It is assumed that the vectors $\{\xi^i : 1 \le i \le n\} \subset \mathbb{R}^\ell$ are linearly independent.

WORKLOAD RELAXATION FOR THE FLUID MODEL
The differential inclusion satisfying the following state-space and rate constraints:

(i) $\widehat{q}(t) \in \mathbb{R}^\ell_+$ for each $t \ge 0$;
(ii) $(t-s)^{-1}(\widehat{q}(t) - \widehat{q}(s)) \in \widehat{\mathsf{V}}$ for each $0 \le s < t$, where

$$\widehat{\mathsf{V}} = \left\{v : \langle \xi^i, v\rangle \ge -\delta_i, 1 \le i \le n\right\}.$$

For a given initial condition $x \in \mathbb{R}^\ell_+$, we say that $\widehat{\boldsymbol{q}}$ is *admissible* if it is a measurable function of time, and satisfies these constraints with $\widehat{q}(0) = x$. An admissible solution need not be continuous since $\widehat{\mathsf{V}}$ is in general unbounded.

Let Ξ denote the $n \times \ell$ matrix whose rows are equal to $\{\xi^i\}$. The *workload process* and *idleness process* are defined respectively by,

$$\widehat{w}(t) := \Xi\widehat{q}(t), \quad I(t) := \widehat{w}(t) - \widehat{w}(0) + t\delta, \qquad t \ge 0.$$

The following properties follow from the definitions (see [16, p. 188]), and demonstrate that $\widehat{\boldsymbol{w}}$ is a generalization of the relaxation (6) for the KSRS model. The bounds in (21) are decoupled since the workload vectors $\{\xi^i : 1 \le i \le n\}$ are assumed linearly independent.

Proposition 3.1: The following hold for each initial condition $x \in \mathbb{R}^\ell_+$, and each admissible trajectory starting from x:

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

$$\frac{d^+}{dt}\widehat{w}_i(t) \ge -\delta_i, \qquad 1 \le i \le n. \tag{21}$$

(ii) The idleness process is *non-decreasing*: Setting $\iota(t) = \frac{d^+}{dt}I(t)$, we have

$$\iota_i(t) \ge 0, \qquad t \in \mathbb{R}_+,\ i = 1,\dots,n. \tag{22}$$

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

$$\widehat{w}(t) \in \mathsf{W} := \{\Xi x : x \in \mathbb{R}^\ell_+\}. \tag{23}$$

□

On constructing control solutions for $\widehat{\boldsymbol{q}}$ we restrict to the workload process, where control amounts to determining the idleness process $\boldsymbol{I}$. For this purpose, suppose that a cost function is given, $c\colon \mathbb{R}^\ell \to \mathbb{R}_+$. We assume that c is piecewise linear, and defines a norm on $\mathbb{R}^\ell$. The associated effective cost $\bar{c}\colon \mathsf{W} \to \mathbb{R}_+$ is defined for $w \in \mathsf{W}$ as the solution to the convex program,

$$\bar{c}(w) = \mathbf{min}\ \ c(x) \qquad \mathbf{s.\,t.}\ \ \Xi x = w,\ x \in \mathbb{R}^\ell_+. \tag{24}$$

Just as was seen in the KSRS model examined in Section II, the effective cost is piecewise linear, of the form

$$\bar{c}(w) = \max_{1 \le i \le n_c} \langle \bar{c}^i, w\rangle, \quad w \in \mathsf{W}, \tag{25}$$

where $\{\bar{c}^i : 1 \le i \le n_c\} \subset \mathbb{R}^n$. The region on which $\bar{c}$ is monotone is denoted,

$$\mathsf{W}^+ = \{w \in \mathsf{W} : \bar{c}(w) \le \bar{c}(w') \quad \text{if } w' \in \mathsf{W},\ w' \ge w\}. \tag{26}$$

We let $\mathcal{X}^*\colon \mathsf{W} \to \mathbb{R}^\ell_+$ denote a continuous function such that $\mathcal{X}^*(w)$ an optimizer of (24) for each $w \in \mathsf{W}$. Then, given a desirable solution $\widehat{\boldsymbol{w}}$ on W, an admissible solution is specified by

$$\widehat{q}(t) = \mathcal{X}^*(\widehat{w}(t)), \qquad t \ge 0.$$

We consider exclusively affine policies to determine the idleness process $\boldsymbol{I}$. For the general workload model, we again fix a polyhedral region $\mathcal{R} \subset \mathsf{W}$ containing the origin with non-empty interior. Under these assumptions, there exist non-negative constants $\{\beta_i\}$ and vectors $\{n_i\} \subset \mathbb{R}^n$ such that,

$$\mathcal{R} = \{w \in \mathbb{R}^n : \langle n^i, w\rangle \ge -\beta_i,\ 1 \le i \le n_R\}. \tag{27}$$

The following additional assumption is assumed so that we can construct a *minimal process* on $\mathcal{R}$ in analogy with the construction in Section II: For each $y \in \mathbb{R}^n$ there exists a 'minimal element' $[y]_+$ satisfying

(i) $[y]_+ \in \mathcal{R}$; (ii) If $y' \in \mathcal{R}$ and $y' \ge y$, then $y' \ge [y]_+$.

The existence of the pointwise projection requires some assumptions on the set $\mathcal{R}$ for dimensions 2 or higher [16].

The minimal process starting from an initial condition $w \in \mathsf{W}$ is then defined by,

$$\widehat{w}(0) = w, \quad \widehat{w}(t) = [w - t\delta]_+, \qquad t > 0. \tag{28}$$

It is represented as the solution to the ODE,

$$\frac{d^+}{dt}\widehat{w}(t) = -\delta + \iota(t), \qquad t > 0,\ \widehat{w}(0+) = [w]_+, \tag{29}$$

where the idleness rate $\boldsymbol{\iota}$ is expressed as the state-feedback law $\iota(t) = \phi(\widehat{w}(t))$, $t \geq 0$, with

$$\begin{aligned} \phi(w) &:= \delta - \delta(w); \\ \delta(w) &:= \lim_{t\downarrow 0} t^{-1}\{[w - t\delta]_+ - w\}, \quad w \in \mathcal{R}. \end{aligned} \tag{30}$$

The mapping ϕ defines the affine policy with respect to the region $\mathcal{R}$.

We conclude with some properties of the projection, and resulting properties of the feedback law ϕ. The ith face $F(i) \subset \mathcal{R}$ is defined by

$$F(i) := \{w \in \mathcal{R} : \langle n^i, w\rangle = -\beta_i\}\,, \qquad 1 \leq i \leq n_R. \tag{31}$$

We assume without loss of generality that each of these sets is of dimension $n-1$.

Proposition 3.2: Suppose that the set $\mathcal{R}$ is given in (27), and that the pointwise projection $[\,\cdot\,]_+ \colon \mathbb{R}^n \to \mathcal{R}$ exists. Then,

(i) For each $y \in \mathbb{R}^n$, the projection $[y]_+$ is the unique optimizer $w^* \in \mathcal{R}$ of the linear program,

$$\begin{array}{lrcl} \textbf{min} & w_1 + \cdots + w_n & & \\ \textbf{s. t.} & \langle n^i, w\rangle & \geq & -\beta_i\,,\ 1 \leq i \leq n_R \\ & w & \geq & y. \end{array}$$

(ii) For each face $F(i)$ with $i \in \{1, \dots n_R\}$, there is a unique $j_i \in \{1, \dots, n\}$ satisfying,

$$n^i_{j_i} > 0\,, \quad \text{and} \quad n^i_j \leq 0\,,\ j \neq j_i.$$

(iii) The feedback law $\phi\colon \mathsf{W} \to \mathbb{R}^n_+$ defined in (31) satisfies the following properties: $\phi(w) = 0$ for $w \in \text{interior}\,(\mathcal{R})$. Otherwise, if for some $i \in \{1, \dots, n_R\}$ we have $w \in F(i)$, and $w \notin F(i')$ for $i' \neq i$, then $\phi(w)_j = 0$ for $j \neq j_i$, where j_i is defined in (ii).

PROOF: Let $w^\circ \in \mathcal{R}$ be an optimizer of the linear program in (i). We have $w^\circ \geq [y]_+$ since $w^\circ \geq y$, and hence also $\sum w^\circ_i \geq \sum([y]_+)_i$. It follows that $[y]_+$ is the unique optimizer of this linear program.

We prove part (ii) by contradiction: Fix $1 \leq i \leq n_R$, and suppose that in fact $n^i_j > 0$ and $n^i_k > 0$ for some $1 \leq j < k \leq n_R$.

Consider $w \in F(i)$, with $w \notin F(i')$ for $i' \neq i$. For $\varepsilon > 0$ we consider the open ball centered at w given by $B(w, \varepsilon) = \{y \in \mathbb{R}^n : \|w - y\| < \varepsilon\}$, where $\|\cdot\|$ denotes the Euclidean norm. Then, from the definition (31), there exists $\varepsilon > 0$ such that whenever $y \in B(w, \varepsilon)$ we have $[y]_+ \notin F(i')$ for $i' \neq i$.

Fix $y \in B(w, \varepsilon)$ with $y \notin \mathcal{R}$, and define

$$s_j = -(n^i_j)^{-1}(\beta_i + \langle n^i, y\rangle), \quad s_k = -(n^i_k)^{-1}(\beta_i + \langle n^i, y\rangle).$$

We have $s_j > 0$, $s_k > 0$ by construction, and the vectors $\{w^j := y + s_j e^j,\ w^k := y + s_k e^k\}$ satisfy $\langle n^i, w^j\rangle = \langle n^i, w^k\rangle = -\beta_i$. Hence, by reducing $\varepsilon > 0$ if necessary, we may assume that each of these vectors lies in $F(i) \subset \mathcal{R}$. We conclude that $[y]_+ \leq \min(w^j, w^k)$, which is only possible if $[y]_+ = y$. This violates our assumptions, and completes the proof of (ii).

Part (iii) is immediate from (ii) and the definition (30). $\square$

B. Brownian workload model

The stochastic workload model considered in the remainder of this paper is defined as follows:

CONTROLLED BROWNIAN-MOTION WORKLOAD MODEL The continuous-time workload process obeys the dynamics,

$$\widehat{W}(t) = w - \delta t + I(t) + N(t), \qquad \widehat{W}(0) = w \in \mathsf{W}\,. \tag{32}$$

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

(i) $\widehat{\boldsymbol{W}}$ is constrained to the workload space W (see (23)).

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

(iii) The idleness process $\boldsymbol{I}$ is adapted to the Brownian motion $\boldsymbol{N}$, with $I(0) = 0$, and $I_i(t) - I_i(s) \geq 0$ for $t \geq s$, $1 \leq i \leq n$.

A stochastic process $\widehat{\boldsymbol{W}}$ satisfying these constraints is called *admissible*.

The definition of an affine policy for the CBM model requires some effort. We begin with the following refinement of the deterministic fluid model defined in Section III-A through the introduction of an additive disturbance: Fix a continuous function $\boldsymbol{d}\colon \mathbb{R}_+ \to \mathbb{R}^n$ satisfying $d(0) = 0$, and consider the controlled model,

$$\widehat{w}(t) = w - t\delta + I(t) + d(t)\,, \qquad t \geq 0. \tag{33}$$

A deterministic process $\widehat{\boldsymbol{w}}$ on W is again called admissible if (33) holds for some non-decreasing idleness process $\boldsymbol{I}$, with $I(0) = 0$.

We say that $\widehat{\boldsymbol{w}}$ is the minimal solution on the region $\mathcal{R}$, with initial condition $w \in \mathsf{W}$, if (i) $\widehat{\boldsymbol{w}}$ is admissible; and (ii) $\widehat{w}'(t) \geq \widehat{w}(t)$ for all $t \geq 0$, and any other admissible solution $\widehat{\boldsymbol{w}}'$ with initial condition w. A minimal solution is a particular instance of a solution to the Skorokhod problem (e.g. [39], [40], [41].) The reflection direction on the ith face is given by e^{j_i}, the j_ith standard basis vector in $\mathbb{R}^n$, where j_i is defined in Proposition 3.2. Under general conditions it is known that the Skorokhod map is a Lipschitz-continuous functional of $\boldsymbol{d}$, and the initial condition w (see [42]).

The following result is a minor extension of [16, Theorem 3.10]. The L_∞-norm in Proposition 3.3 (ii) is given by,

$$\|\boldsymbol{d} - \boldsymbol{d}'\|_{[0,t]} := \max_{0 \leq s \leq t} \|d(s) - d'(s)\|, \qquad t \geq 0,$$

where $\|\cdot\|$ denotes the Euclidean norm.

Proposition 3.3: Suppose that the set $\mathcal{R}$ is given in (27), and that the pointwise projection $[\,\cdot\,]_+ \colon \mathbb{R}^n \to \mathcal{R}$ exists. Then,

(i) The minimal process $\widehat{\boldsymbol{w}}$ exists for each continuous disturbance $\boldsymbol{d}$, and each initial condition $w \in \mathsf{W}$.

(ii) There exists a fixed constant $k_R < \infty$ satisfying the following: For any two initial conditions $w, w' \in \mathsf{W}$, and any two continuous disturbance processes $\boldsymbol{d}, \boldsymbol{d}'$, the resulting workload processes $\{\widehat{w}(\,\cdot\,; w), \widehat{w}'(\,\cdot\,; w')\}$ obtained with the respective initial conditions and disturbances satisfy the following bounds for $t \geq 0$:

$$\|\widehat{w}(t; w) - \widehat{w}'(t; w')\| \leq k_R\big[\|w - w'\| + \|\boldsymbol{d} - \boldsymbol{d}'\|_{[0,t]}\big].$$

(iii) For each $\widehat{w}(0) \in \mathsf{W}$, $i \in \{1, \dots, n_R\}$, and $j \neq j_i$,

$$\int_0^\infty \mathbb{I}\big(\widehat{w}(t) \in F(i),\ \widehat{w}(t) \notin F(i') \text{ for } i' \neq i\big)\, dI_j(t) = 0.$$

□

Proposition 3.3 leads to the following definition of the affine policy for (32) with respect to a domain $\mathcal{R}$:

$$(\widehat{\boldsymbol{W}}, \boldsymbol{I}) = \textit{the minimal process on } \mathcal{R} \textit{ with } \boldsymbol{d} = \boldsymbol{N}. \tag{34}$$

The resulting idleness process $\boldsymbol{I}$ is random since $\boldsymbol{N}$ is assumed to be a Gaussian stochastic process. The controlled process $\widehat{\boldsymbol{W}}$ is a time-homogeneous, strong Markov process. The strong Markov property follows from the sample-path construction of $\widehat{\boldsymbol{W}}$ [40].

It will be convenient below to introduce the *extended generator* $\mathcal{A}$ for the Markov process $\widehat{\boldsymbol{W}}$ under an affine policy. A stochastic process $\boldsymbol{M}$ adapted to some filtration $\{\mathcal{F}_t : t \geq 0\}$ is called a *martingale* if $\mathsf{E}[M(t{+}s) \mid \mathcal{F}_t] = M(t)$ for each $t, s \in \mathbb{R}_+$. It is called a *local-martingale* if there exists an increasing sequence of stopping times $\{\varsigma_n\}$ such that $\{M(t \wedge \varsigma_n) : t \in \mathbb{R}_+\}$ is a martingale, for each $n \geq 1$, and $\varsigma_n \uparrow \infty$ a.s. as $n \to \infty$. Throughout the paper we assume that $\{\mathcal{F}_t : t \geq 0\}$ is the filtration generated by $\boldsymbol{N}$.

A measurable function $f \colon \mathcal{R} \to \mathbb{R}$ is in the domain of $\mathcal{A}$ if there is a measurable function $g \colon \mathcal{R} \to \mathbb{R}$ such that, for each initial condition $\widehat{W}(0) \in \mathcal{R}$, the stochastic process defined below for $t \geq 0$ is a local-martingale,

$$M_f(t) := f(\widehat{W}(t)) - \big\{ f(\widehat{W}(0)) + \int_0^t g(\widehat{W}(s))\, ds \big\}. \tag{35}$$

In this case, we write $g = \mathcal{A}f$.

The following version of Itô's formula (37) is used to identify a large class of functions within the domain of $\mathcal{A}$. This result will be applied repeatedly in the treatment of value functions that follows. Given a function $f \colon \mathsf{W} \to \mathbb{R}$, the differential generator is expressed

$$\mathcal{D}f := -\delta^{\mathsf{T}} \nabla f + \tfrac{1}{2} \Delta f\,, \tag{36}$$

where $\Delta f := \sum_{i,j} \big(\Sigma(i,j) \frac{\partial^2}{\partial_{x_i} \partial_{x_j}} f\big)$ is the weighted Laplacian.

Theorem 3.4: Suppose that the set $\mathcal{R}$ is given in (27), and that the pointwise projection $[\,\cdot\,]_+ \colon \mathbb{R}^n \to \mathcal{R}$ exists. Suppose that $f \colon \mathbb{R}^n \to \mathbb{R}$ is continuously differentiable, ∇f is uniformly Lipschitz continuous on compact sets, and that Δf exists for a.e. $w \in \mathcal{R}$. Then,

(i) For each initial condition $\widehat{W}(0) \in \mathcal{R}$,

$$\begin{aligned} f(\widehat{W}(t)) = f(\widehat{W}(0)) &+ \int_0^t [\mathcal{D}f](\widehat{W}(s))\, ds \\ &+ \int_0^t \langle \nabla f(\widehat{W}(s)), dI(s) \rangle \\ &+ \int_0^t \langle \nabla f(\widehat{W}(s)), dN(s) \rangle \end{aligned} \tag{37}$$

(ii) Suppose that $\delta \in \operatorname{interior}(\mathcal{R})$, and in addition the following bounary conditions hold,

$$\langle \phi(w), \nabla f(w) \rangle = 0\,, \qquad w \in \mathcal{R}\,, \tag{38}$$

where ϕ is defined in (30). Then f is in the domain of $\mathcal{A}$, and $\mathcal{A}f = \mathcal{D}f$.

(iii) If the conditions of (ii) hold, and in addition,

$$\mathsf{E}\Big[\int_0^t \|\nabla f(\widehat{W}(s;w))\|^2\, ds\Big] < \infty, \qquad w \in \mathcal{R},\ t \geq 0,$$

then $\boldsymbol{M}_f$ is a martingale for each initial condition.

PROOF: When f is C^2, then (i) is given in [43, Theorem 2.9, p. 287]. In the more general setting given in the theorem, the function f may be approximated uniformly on compacta by C^2 functions so that $\mathcal{D}f$ is simultaneously approximated in $L_2(C)$ for any compact set $C \subset \mathcal{R}$. See [44, Theorem 1, p. 122], and the extensions in [45], [46], [47].

Suppose now that the assumptions of (ii) hold. The assumption $\delta \in \operatorname{interior}(\mathcal{R})$ is imposed to ensure that ϕ does not vanish on $\partial\mathcal{R}$. Itô's formula then gives the following representation,

$$M_f(t) = \int_0^t \langle \nabla f(\widehat{W}(s)), dN(s) \rangle, \qquad t \geq 0. \tag{39}$$

The local martingale property is immediate since $\boldsymbol{N}$ is driftless Brownian motion.

We have for each $0 \leq s < t$, and each initial w,

$$\mathsf{E}[(M_f(t) - M_f(s))^2] = \int_s^t \mathsf{E}[(\nabla f(\widehat{W}(s)))^{\mathsf{T}} \Sigma \nabla f(\widehat{W}(s))].$$

The right hand side is finite under the conditions of (iii), and the martingale property then follows from the representation (39) [43]. □

IV. VALUE FUNCTIONS

We now develop structural properties of the value functions associated with the fluid and CBM models. At the same time, we establish a strong form of geometric ergodicity for the process $\widehat{\boldsymbol{W}}$ under an affine policy.

It will be useful to introduce a scaling parameter $\kappa \geq 0$ to investigate the impact of variability,

$$\widehat{W}(t) = w - \delta t + I(t) + \sqrt{\kappa} N(t), \qquad \widehat{W}(0) = w \in \mathsf{W}\,, \tag{40}$$

in which the idleness process is determined by the affine policy (34). We restrict to affine domains of the form,

$$\mathcal{R}(\kappa) = \{ w \in \mathbb{R}^n_+ : \langle n^i, w \rangle \geq -\kappa \beta_i\,,\ 1 \leq i \leq n_R \}, \tag{41}$$

where $\kappa \geq 0$ is the constant used in (40), $\beta \in \mathbb{R}^{n_R}_+$ is a constant vector, and $\{n^i\} \subset \mathbb{R}^n$. This is motivated by affine approximations of optimal policies of the form obtained in [17] and reviewed in Section II-C.

We consider two processes on the domain $\mathcal{R}(\kappa)$: The minimal process $\widehat{\boldsymbol{W}}$ satisfying (40); and also the deterministic minimal solution $\widehat{\boldsymbol{w}}$ on $\mathcal{R}(\kappa)$ defined in (28). When we wish to emphasize the dependency on κ and the initial condition $w \in \mathsf{W}$ we denote the workload processes by $\widehat{W}(t; w, \kappa)$, $\widehat{w}(t; w, \kappa)$, respectively.

Let $\widehat{\eta}_\kappa$ denote the steady-state mean of $\bar{c}(\widehat{W}(t; w, \kappa))$. When $\kappa = 1$ we drop the subscript so that $\widehat{\eta} = \widehat{\eta}_1$. The value functions considered in this section are defined for the fluid and CBM workload models respectively by,

$$J(w;\kappa) \;:=\; \int_0^\infty \bar c(\widehat w(t;w,\kappa))\,dt\,; \tag{42}$$

$$h(w;\kappa) \;:=\; \int_0^\infty \Big(\mathsf{E}[\bar c(\widehat W(t;w,\kappa))] - \widehat\eta_\kappa\Big)\,dt, \tag{43}$$

where $w \in \mathsf{W}$ and $\kappa \geq 0$. We again suppress dependency on κ when $\kappa = 1$.

Consider for the moment the special case $\kappa = 1$. Assuming that the integral in (43) exists and is finite for each w, we obtain from the Markov property the following representation for each $T > 0$,

$$h(\widehat W(T;w)) = \int_T^\infty \Big(\mathsf{E}[\bar c(\widehat W(t;w)) \mid \mathcal{F}_T] - \widehat\eta\Big)\,dt.$$

It then follows that the stochastic process $\boldsymbol{M}_h$ defined below is a martingale for each initial condition $w \in \mathsf{W}$,

$$\begin{aligned} M_h(t) &:= h(\widehat W(t;w)) - h(w) \\ &\quad + \int_0^t \big[\bar c(\widehat W(s;w)) - \widehat\eta\big]\,ds, \quad t \geq 0. \end{aligned} \tag{44}$$

Hence h is in the domain of the extended generator, and solves the dynamic programming equation,

$$\mathcal{A}h = -\bar c + \widehat\eta. \tag{45}$$

This is also called *Poisson's equation*, and the function h is known as the *relative value function* [48]. General conditions under which both the steady-state mean $\widehat\eta$ and the integral in (43) are well defined are presented in Theorem 4.5.

We begin in Section IV-A with the simpler fluid model. The proof of Proposition 2.1 can be generalized to conclude that the value function J is continuous and piecewise-quadratic in the general fluid model considered here (see e.g. [49, Section 4.2].) Our goal is to show that J is a smooth solution to the dynamic programming equation (12), which is equivalently expressed,

$$\mathcal{D}_0 J = -\bar c \tag{46}$$

where the differential generator for the fluid model is defined in analogy with (36) via,

$$\mathcal{D}_0 f := -\delta^{\scriptscriptstyle T}\nabla f, \qquad f \in C^1(\mathsf{W}). \tag{47}$$

To show that J is smooth we differentiate the workload process $\widehat w(t) = [w - \delta t]_+$ with respect to the initial condition w. Let $\mathcal{O} \subset \mathbb{R}^n$ denote the maximal open set such that $[y]_+$ and $\bar c([y]_+)$ are each C^1 for $y \in \mathcal{O}$. Generally, since $\bar c$ and the projection are each piecewise linear on $\mathbb{R}^n$, it follows that the set $\mathcal{O}$ can be expressed as the union,

$$\mathcal{O} = \bigcup_{1\leq i\leq n_O} R_i, \qquad \bar{\mathcal{O}} = \mathbb{R}^n. \tag{48}$$

where each of the sets $\{R_i : 1 \leq i \leq n_O\}$ is an open polyhedron. Note that $\mathcal{O}$ and the sets $\{R_i\}$ will in general depend upon κ.

For example, consideration of the KSRS model with $\mathcal{R} = \mathbb{R}^2_+$, we see from Figure 2 that $[\,\cdot\,]_+$ and $\bar c([\,\cdot\,]_+)$ are both linear on each of the sets $\{R_1, R_2, R_3\}$ shown in Figure 3. The open polyhedron $R_4 = \{y \in \mathbb{R}^2 : y < 0\}$ is also contained in $\mathcal{O}$ since $[y]_+ = 0$ and $\bar c([y]_+) = 0$ on R_4. This reasoning leads to a representation of the form (48) when $\mathcal{R} = \mathbb{R}^2_+$.

We list here some key assumptions imposed in the results that follow:

(A1) For each $\kappa \geq 0$, the set $\mathcal{R}(\kappa)$ has non-empty interior, satisfies $\mathcal{R}(\kappa) \subset \mathbb{R}^n_+$, and the pointwise projection $[\,\cdot\,]_+ \colon \mathbb{R}^n \to \mathcal{R}(\kappa)$ exists.

(A2) $\displaystyle\int_0^\infty \mathbb{I}\{(w - \delta t) \in \mathcal{O}^c\}\,dt = 0$ for each $w \in \mathbb{R}^n$, $\kappa \geq 0$.

(A3) $\delta \in \text{interior}\,(\mathcal{R}(0))$.

A. The fluid value function

In our consideration of the fluid value function we require some structural properties for $\widehat{\boldsymbol{w}}$. The following scaling property follows from the definitions:

Lemma 4.1: $\widehat w(t;w,\kappa) = \kappa\widehat w(\kappa^{-1}t;\kappa^{-1}w,1)$ for each $t \geq 0$, $\kappa > 0$, and $w \in \mathcal{R}(\kappa)$. □

Theorem 4.2 establishes several useful properties of the fluid value function. To illustrate the conclusions and assumptions we again turn to the KSRS model. Consider first the case in which $\delta_1 = \delta_2 > 0$, the cost function is $\bar c(w) = \max(w_1, w_2)$, and the constraint region $\mathcal{R} = \mathsf{W} = \mathbb{R}^2_+$. Assumption (A2) does not hold: The integral in (A2) is non-zero for any non-zero initial condition w on the diagonal in $\mathbb{R}^2_+$. The value function is given by $J(w) = \frac{1}{2}\delta_1^{-1}\max(w_1^2, w_2^2)$, which is not C^1 on $\mathbb{R}^2_+$. This explains why (A2) is needed in Theorem 4.2 (i).

To see why (A3) is required in Theorem 4.2 (iii) we take $\delta_1 = 4\delta_2$; $\bar c(w) = w_1 + w_2$; and $\mathcal{R} = \{0 \leq w_1 \leq 3w_2 \leq 9w_1\}$. Assumptions (A1) and (A2) hold, and the C^1 value function may be explicitly computed:

$$J(w) = \tfrac{1}{2} w^{\scriptscriptstyle T} D w,\ w \in \mathcal{R}, \quad \text{with } D = \tfrac{1}{11}\delta_1^{-1}\begin{bmatrix} 3 & -1 \\ -1 & 15 \end{bmatrix}.$$

Although smooth, we have $\frac{\partial}{\partial w_2} J(w) \neq 0$ and $\phi(w) = 0$ along the lower boundary of $\mathcal{R}$. This is possible since the model violates (A3).

We note that the value function for the discrete-time constrained linear quadratic regulator problem is also smooth and piecewise quadratic under general conditions [50], [51].

Theorem 4.2: Under (A1) we have $J(w;\kappa) < \infty$ for each $w \in \mathsf{W}$, $\kappa > 0$, and the following scaling property holds:

$$J(w;\kappa) = \kappa^2 J(\kappa^{-1}w;1)\,, \qquad w \in \mathcal{R}(\kappa)\,,\ \kappa > 0. \tag{49}$$

If in addition (A2) holds, then for each $\kappa \geq 0$,

(i) The function $J \colon \mathcal{R}(\kappa) \to \mathbb{R}_+$ is piecewise-quadratic, C^1, and its gradient ∇J is globally Lipschitz continuous on $\mathcal{R}(\kappa)$.

(ii) The dynamic programming equation (46) holds on $\mathcal{R}(\kappa)$.

(iii) The boundary conditions (38) hold for J, where $\phi(w)$ is defined in (30) with respect to the region $\mathcal{R}(\kappa)$. If in addition (A3) holds, then the function ϕ does not vanish on $\partial\mathcal{R}(\kappa)$.

PROOF: The scaling property (49) follows from Lemma 4.1.

The dynamic programming equation (ii) is simply the fundamental theorem of calculus: See (11) and surrounding discussion.

To establish the smoothness property in (i) we define for $t \geq 0$, $\kappa \geq 0$, $w \in \mathcal{R}(\kappa)$,

$$\Gamma(t; w, \kappa) := \nabla_w \{\bar{c}(\widehat{w}(t; w, \kappa))\}, \tag{50}$$

whenever the gradient with respect to w exists. Let $\Pi\colon \mathcal{O} \to \mathbb{R}^{n\times n}$ denote the $n \times n$ derivative of $[\,\cdot\,]_+$ with respect to w. The matrix-valued function Π is constant on each of the connected components of $\mathcal{O}$. The gradient (50) exists whenever $y(t) := w - \delta t \in \mathcal{O}$ for a given $w \in \mathcal{R}$, $t > 0$, and by the chain rule,

$$\Gamma(t; w, \kappa) = \Pi(y(t))^{\scriptscriptstyle T} \big(\{\nabla_w \bar{c}\}\,([y(t)]_+)\big).$$

Under (A2) we have $y(t) \in \mathcal{O}$ for each $w \in \mathcal{R}(\kappa)$, and a.e. $t \in \mathbb{R}_+$.

We conclude that, for each $w \in \mathsf{W}$, the gradient (50) exists for a.e. $t \geq 0$. Lipschitz continuity of $\bar{c}(\widehat{w}(t; w, \kappa))$ in (t, w) then leads to the representation,

$$\nabla_w J(w; \kappa) = \int_0^{\widehat{T}^*(w)} \Gamma(t; w, \kappa)\, dt, \tag{51}$$

where the minimal draining time for the fluid model is the piecewise linear function,

$$\widehat{T}^*(w) = \max_{1\leq i\leq n} \delta_i^{-1} w_i, \qquad w \in \mathsf{W}.$$

The range of integration in (51) is finite since $\Gamma(t; w, \kappa) = 0$ for $t > \widehat{T}^*(w)$.

The proof of (i) is completed on establishing Lipschitz continuity of the right hand side of (51). For a given initial condition $w \in \mathcal{R}(\kappa)$, let $\{R_i^w : i = 1, \dots, m\}$ denote the sequence of regions in $\{R_i\}$ (defined in (48)) so that $y(t) := w - \delta t$ enters regions $\{R_1^w, R_2^w, \dots\}$ sequentially. We set $T_0 = 0$, let T_i denote the exit time from R_i^w, and set $\Gamma_i := \Gamma(t; w, \kappa)$ for $t \in (T_{i-1}, T_i)$, $i \geq 1$. From the forgoing we see that the gradient can be expressed as the finite sum,

$$\nabla_w J(w; \kappa) = \sum_{i=1}^{m-1} (T_i(w) - T_{i-1}(w))\Gamma_i.$$

Each of the functions $\{T_i(w)\}$ is piecewise linear and continuous on $\mathcal{R}(\kappa)$ under (A1) and (A2), and this implies that ∇J is Lipschitz continuous.

To see (iii), fix $w \in \partial\mathcal{R}(\kappa)$, and $w' \in \text{interior}\,(\mathcal{R}(\kappa))$. Then $w^\theta := (1-\theta)w + \theta w' \in \text{interior}\,(\mathcal{R}(\kappa))$ for each $\theta \in (0, 1]$. Consequently, applying (ii) we have for $\theta \in (0, 1]$,

$$\langle \delta, \nabla J(w^\theta)\rangle = \bar{c}(w^\theta), \textit{ and } \langle \delta(w), \nabla J(w)\rangle = \bar{c}(w).$$

This combined with continuity of ∇J and $\bar{c}$ establishes the boundary property $\langle \phi(w), \nabla J(w)\rangle = 0$. It is obvious from (30) that $\phi(w) \neq 0$ on the boundary under (A3). ☐

B. Regeneration and ergodicity

A stationary version of the workload process can be constructed using the shift-coupling technique [52]. Suppose that the Brownian motion $\boldsymbol{N}$ is defined on the two-sided interval $\mathbb{R}$, with $N(0) = 0$, and construct a process $\widehat{\boldsymbol{W}}^s$ on $\mathcal{R}$, *initialized at time* $-s$. For a given initial condition $w \in \mathcal{R}$, this is defined on the interval $[-s, \infty)$ with initial condition $\widehat{W}^s(-s; w) = w$, and noise process $N^s(t) := N(t) - N(-s)$, $t \geq -s$. Suppose that all of the processes are initialized at $w = 0$. Then, for any fixed t, it can be shown that $\{\widehat{W}^s(t; 0) : s \geq 0\}$ is non-decreasing in s for $s \geq -t$. Consequently, the following limit exists with probability one,

$$\widehat{W}^\infty(t) := \lim_{s\to\infty} \widehat{W}^s(t; 0), \qquad -\infty < t < \infty. \tag{52}$$

It follows from the main result of [24] that the controlled process $\widehat{\boldsymbol{W}}$ is ergodic, so that $\widehat{W}^\infty(t)$ is a.s. finite.

The following generalization of Lemma 4.1 allows us to restrict to $\kappa = 1$. Its proof is immediate from the scaling formula $\kappa N(\kappa^{-1}t) \stackrel{\rm d}{=} \sqrt{\kappa} N(t)$, $t \geq 0$.

Lemma 4.3: $\widehat{W}(t; w, \kappa) \stackrel{\rm d}{=} \kappa \widehat{W}(\kappa^{-1}t; \kappa^{-1}w, 1)$ for each $t \geq 0$, $\kappa > 0$, and $w \in \mathcal{R}(\kappa)$. ☐

As one consequence of Lemma 4.3 we see that the steady-state cost scales linearly, $\eta_\kappa = \kappa\eta_1$ for $\kappa > 0$. Lemma 4.3 combined with Proposition 3.3 also implies a weak form of stability:

Proposition 4.4: If (A1) holds then the minimal process on $\mathcal{R}$ with $\kappa = 1$ satisfies, for each $p \geq 1$, $w \in \mathsf{W}$,

$$\lim_{r\to\infty} \mathsf{E}[\|r^{-1}\widehat{W}(rt; rw)\|^p] = 0, \quad t \geq \widehat{T}^*(w). \tag{53}$$

PROOF: From Proposition 3.3 we obtain the bound, for each $t \geq 0$, $\kappa > 0$, $w \in \mathsf{W}$,

$$\|\widehat{W}(t; w, \kappa) - \widehat{w}(t; w, \kappa)\| \leq k_R \sqrt{\kappa} \|\boldsymbol{N}\|_{[0,t]}. \tag{54}$$

If $t \geq \widehat{T}^*(w)$ then $\widehat{w}(t; w, \kappa) = 0$ for each $\kappa > 0$. Combining the bound (54) with Lemma 4.3 we conclude that, for any $p \geq 1$, $t \geq \widehat{T}^*(w)$,

$$\kappa^p \mathsf{E}[\|\widehat{W}(\kappa^{-1}t; \kappa^{-1}w)\|^p] \leq (k_R\sqrt{\kappa})^p \mathsf{E}[\|\boldsymbol{N}\|_{[0,t]}^p].$$

An application of [53, Corollary 37.12] shows that $\mathsf{E}[\|\boldsymbol{N}\|_{[0,t]}^p]$ is finite for each $t \geq 0$, $p \geq 1$. The conclusion (53) is obtained on setting $r = \kappa^{-1}$. ☐

The limit (53) is precisely the form of fluid-scale stability assumed in [27] to formulate criteria for moments in a stochastic network. For a given function $f\colon \mathcal{R} \to [1, \infty)$, and for a pair of probability distributions μ, ν on $\mathcal{B}(\mathcal{R})$ we define,

$$\|\mu - \nu\|_f := \sup_{|g|\leq f} |\mu(g) - \nu(g)|.$$

Define $f_p(w) = \|w\|^p + 1$ for $w \in \mathsf{W}$, and let $\{P^t : t \geq 0\}$ denote the Markov transition group for $\widehat{\boldsymbol{W}}$ with $\kappa = 1$.

The main result of [23], [24] establishes positive recurrence of the reflected diffusion $\widehat{\boldsymbol{W}}$. The following result extends these results by establishing polynomial moments, and a polynomial rate of convergence of the underlying distributions.

Theorem 4.5: Suppose that (A1) holds. Then, the minimal process on $\mathcal{R}$ with $\kappa = 1$ satisfies

(i) The limiting process $\widehat{\boldsymbol{W}}^{\infty}$ exists, and its marginal distribution π on $\mathcal{R}$ is the unique steady-state distribution for $\widehat{\boldsymbol{W}}$.

(ii) The steady-state distribution has finite moments, and $\lim_{t\to\infty} t^p \|P^t(w,\,\cdot\,)-\pi(\,\cdot\,)\|_{f_p} = 0$ for each $p \ge 1$, $w \in \mathsf{W}$.

(iii) There is a compact set $C_0 \subset \mathsf{W}$ s. t. for each integer $p \ge 0$, there is a finite constant k_p satisfying,

$$\mathsf{E}\Big[\int_0^{\tau_{C_0}} f_p(\widehat{W}(t;w))\,dt\Big] \le k_p f_{p+1}(w), \quad w \in \mathsf{W}, \tag{55}$$

where τ_{C_0} denotes the first entrance time to C_0.

PROOF: Analogs of (i) and (ii) are obtained in [27, Theorem 6.3]. Although this paper concerns networks with renewal input processes, the proof in the present setting is identical. Part (iii) is [27, Proposition 5.3]. □

Theorem 4.5 can be strengthened under (A1)–(A3) since we may then construct a Lyapunov function based on the fluid value function.

Proposition 4.6: Suppose that (A1)–(A3) hold. Then, with $\widehat{\boldsymbol{W}}$ equal to the minimal process on $\mathcal{R}$ with $\kappa = 1$,

(i) The fluid value function J is in the domain of the extended generator, and $\mathcal{A}J = \mathcal{D}J = -\bar{c} + b_{\text{CBM}}$, where b_{CBM} is defined a.e. on $\mathcal{R}$ by,

$$b_{\text{CBM}}(w) := \tfrac{1}{2}\Delta J\,(w)\,. \tag{56}$$

(ii) The function $V = \sqrt{1+J}$ is in the domain of the extended generator, and for some $\varepsilon_0 > 0$, $k_0 < \infty$, and a compact set $C_0 \subset \mathcal{R}$,

$$\mathcal{A}V = \mathcal{D}V \le -\varepsilon_0 + k_0 \mathbb{I}_{C_0}$$

(iii) The function $V_\vartheta = e^{\vartheta V}$ is in the domain of the extended generator for each $\vartheta > 0$. There exists $\vartheta_0 > 0$ such that the following bound holds for each $\vartheta \in (0, \vartheta_0]$: For finite constants $\varepsilon_\vartheta > 0$, $b_\vartheta > 0$,

$$\mathcal{A}V_\vartheta = \mathcal{D}V_\vartheta \le -\varepsilon_\vartheta V_\vartheta + b_\vartheta\,.$$

PROOF: Theorem 4.2 (iii) implies that J satisfies the boundary conditions required in Theorem 3.4 (ii), and hence J is in the domain of $\mathcal{A}$. Each of the functions considered in (ii) and (iii) also satisfies the conditions of Theorem 3.4 (ii) since these properties are inherited from J.

The bounds on $\mathcal{D}$ in (ii) and (iii) follow from the identity $\mathcal{D}J = -\bar{c} + b_{\text{CBM}}$, and straightforward calculus. □

We say that $\widehat{\boldsymbol{W}}$ is V_ϑ*-uniformly ergodic* if for some $d_\vartheta > 0$, $b_\vartheta < \infty$, and all $t \ge 0$, $w \in \mathsf{W}$,

$$\|P^t(w,\,\cdot\,) - \pi(\,\cdot\,)\|_{V_\vartheta} \le b_\vartheta V_\vartheta(w) e^{-d_\vartheta t}\,.$$

See [26], [3], [54], [55] for background.

Theorem 4.7: Suppose that (A1)–(A3) hold with $\kappa = 1$. Then, the steady-state mean satisfies $\widehat{\eta} = \pi(\bar{c}) = \pi(b_{\text{CBM}})$, and the minimal process $\widehat{\boldsymbol{W}}$ is V_ϑ-uniformly ergodic for each $\theta \in (0, \vartheta_0]$.

PROOF: V_ϑ-uniform ergodicity follows from [25, Theorem 6.1] and Proposition 4.6 (iii).

The identity $\widehat{\eta} = \pi(b_{\text{CBM}})$ is obtained as follows. We have from Theorem 3.4 the representation,

$$J(\widehat{W}(t)) = J(\widehat{W}(0)) + \int_0^t \big(-c(\widehat{W}(s)) + b_{\text{CBM}}(\widehat{W}(s))\big)\,ds + \int_0^t \langle \nabla J(\widehat{W}(s)), dN(s)\rangle.$$

Setting $\widehat{W}(0) \sim \pi$ and taking expectations gives,

$$t\pi(\bar{c}) = \mathsf{E}\Big[\int_0^t \bar{c}(\widehat{W}(s))\Big] = \mathsf{E}\Big[\int_0^t b_{\text{CBM}}(\widehat{W}(s))\Big] = t\pi(b_{\text{CBM}}).$$

□

C. The relative value function

The main result of this section is Theorem 4.9, which provides a bound on the error $\mathcal{E}(w) = h(w;1) - J(w;1)$ between the relative value function and the fluid value function. We first obtain bounds on h based on the dynamic programming equation (46), along with an analog of Theorem 4.2.

Theorem 4.8: Under (A1) we have $h(w;\kappa) < \infty$ for all $w \in \mathsf{W}$, $\kappa \ge 0$, and the following scaling property holds:

$$h(w;\kappa) = \kappa^2 h(\kappa^{-1}w;1)\,, \qquad w \in \mathcal{R}(\kappa)\,,\ \kappa > 0. \tag{57}$$

Moreover, for each $\kappa > 0$,

(i) For some constant $k > 0$, and all $w \in \mathcal{R}(\kappa)$,

$$-k\kappa^2 \le h(w;\kappa) \le k(\kappa^2 + \|w\|^2), \qquad w \in \mathsf{W}.$$

(ii) If h' is another solution to Poisson's equation (45) that is bounded from below, then there is a constant k' such that $h'(w;\kappa) = h(w;\kappa) + k'$ on $\mathcal{R}(\kappa)$.

PROOF: The scaling property (57) follows directly from Lemma 4.3 and the formula (43). Without loss of generality we restrict to $\kappa = 1$ in the remainder of the proof.

To prove (i) note first that Theorem 4.5 (iii) implies that $\widehat{\boldsymbol{W}}$ is 'f_p*-regular, with bounding function* f_{p+1}' [26], [56], where $f_p(w) = \|w\|^p + 1$ as above. This means that for each p there exists $k_p < \infty$ such that for any set S with positive π-measure there is $k_S < \infty$ satisfying for all w,

$$\mathsf{E}\Big[\int_0^{\tau_S} f_p(\widehat{W}(s;w))\,ds\Big] \le k_p f_{p+1}(w) + k_S. \tag{58}$$

Specializing to $p = 1$, it then follows from [48, Theorem 3.2] that a solution g to Poisson's equation exists with quadratic growth satisfying for each $T \ge 0$,

$$\mathsf{E}[g(\widehat{W}(T;w))] = g(w) - \int_0^T \Big(\mathsf{E}[\bar{c}(\widehat{W}(t;w))] - \widehat{\eta}\Big)\,dt.$$

Letting $T \to \infty$ and applying Theorem 4.5 (ii) then gives $\pi(g) = g(w) - h(w)$, with h defined in (43), so that $|h|$ is also bounded by a quadratic function of w.

Let $S \subset \mathcal{R}$ be any compact set with non-empty interior. The optional sampling theorem implies that $\{M_h(t \wedge \tau_S) : t \ge 0\}$ is a martingale, and it can be shown using Theorem 4.5 and (58) that it is uniformly integrable. Consequently, we obtain the expression, $\mathsf{E}[M_h(\tau_S)] =$

$$\mathsf{E}\Big[h(\widehat{W}(\tau_S;w)) - h(w) + \int_0^{\tau_S} \bar{c}(\widehat{W}(s;w)) - \widehat{\eta}\,ds\Big] = 0. \tag{59}$$

On setting $S = \{w : \bar{c}(w) \le \widehat{\eta}\}$ we obtain the lower bound, $h(w) \ge \inf_{w' \in S} h(w')$, completing the proof of (i).

The uniqueness result (ii) is given in [6, Theorem A3]. Although stated in discrete time, Section 6 of [6] contains a roadmap explaining how to translate to continuous time. □

Note that the identity (57) implies that $r^{-2}h(rw;1) = h(w;r^{-1}) \to h(w;0)$. Similarly, (49) implies that $r^{-2}J(rw;1) = J(w;r^{-1}) \to J(w;0)$, as $r \to \infty$, and it follows that $r^{-2}|\mathcal{E}(rw)| \to 0$ as $r \to \infty$ since $h(w;0) = J(w;0)$. This is a version of the error bounds explored in [6], [32] (see also [8].) The following result provides a substantial strengthening of this asymptotic bound:

Theorem 4.9: Under (A1)–(A3) the value functions h and J differ by a function with linear growth,

$$\sup_{w \in \mathcal{R}} \frac{|\mathcal{E}(w)|}{1 + \|w\|} < \infty.$$

PROOF: We have noted that $\boldsymbol{M}_h$ is a martingale, and $\boldsymbol{M}_J$ is a martingale by Theorem 3.4 (iii). Consequently, $\boldsymbol{M}_{\mathcal{E}} := \boldsymbol{M}_h - \boldsymbol{M}_J$ is a martingale, and can be expressed for $t \ge 0$ by $M_{\mathcal{E}}(t) =$

$$\mathcal{E}(\widehat{W}(t;w)) - \mathcal{E}(\widehat{W}(0;w)) + \int_0^t \bigl[b_{\text{CBM}}(\widehat{W}(s;w)) - \widehat{\eta}\bigr]\, ds\,. \quad (60)$$

Let $C_0 \subset \mathsf{W}$ denote the compact set found in Theorem 4.5 (iii). The martingale property for $\boldsymbol{M}_h$ implies the representation (59), and combining (55) with Itô's formula (37) we obtain the analogous expression for $\boldsymbol{M}_J$. On subtracting, we obtain for each initial condition,

$$\mathcal{E}(w) = \mathsf{E}\Bigl[\mathcal{E}(\widehat{W}(\tau_{C_0};w)) + \int_0^{\tau_{C_0}} b_{\text{CBM}}(\widehat{W}(s;w)) - \widehat{\eta}\, ds\Bigr] \quad (61)$$

To complete the proof, observe that the function b_{CBM} is bounded and $\mathsf{E}_w[\tau_{C_0}]$ has linear growth by Theorem 4.5 (iii) with $p = 0$. □

An analysis of the differential generator $\mathcal{D}$ reveals that the invariant distribution π is of a computable, exponential form under certain conditions on the CBM model and the polyhedral region $\mathcal{R}$. Unfortunately, these conditions are fragile (see [57], [58] and the references therein.) Computational methods are investigated in [59].

Given the complexity of computation of the invariant distribution, we may turn to simulation in order to evaluate a given policy.

V. SIMULATION

In this section we describe how Theorem 4.2 provides theoretical justification for the construction of a smoothed estimator of the form described in Section II-D.

It is unlikely that one would be interested in simulating a CBM model. However, we expect that the bounds obtained here will have analogs in a discrete queueing model.

The following definitions are all generalizations of those used in our treatment of the KSRS model. The standard estimator is again defined as the sample-path average,

$$\widehat{\eta}(T) = \frac{1}{T}\int_0^T \bar{c}(\widehat{W}(t)))\, dt\,, \quad T > 0. \quad (62)$$

The time-average variance constant (TAVC) is given by

$$\gamma^2 = \lim_{T\to\infty} \frac{1}{T}\mathsf{E}_\pi\Bigl[\Bigl(\int_0^T \bar{c}(\widehat{W}(t)) - \widehat{\eta}\, dt\Bigr)^2\Bigr]. \quad (63)$$

The limit in (63) exists and is finite under the assumptions of Theorem 4.7, and has the alternative representation in terms of the mean quadratic-variation of the martingale $\boldsymbol{M}_h$:

$$\gamma^2 = \frac{1}{T}\mathsf{E}_\pi\bigl[\bigl(M_h(T)\bigr)^2\bigr], \quad T > 0. \quad (64)$$

This follows from [26, Theorem 17.5.3], which establishes a version of (63) for any discrete-time geometrically-ergodic Markov chain. The proof in continuous time is identical [54].

To investigate the impact of variability and load we consider a second scaling of the CBM model,

$$\widehat{W}(t) = w - (1-\varrho)\delta t + I(t) + \sqrt{\kappa}N(t), \quad t \ge 0. \quad (65)$$

The idleness process is again determined by the affine policy using the region $\mathcal{R}(\kappa)$ defined in (41). The parameter $0 \le \varrho < 1$ is interpreted as *network load*, and $\kappa > 0$ again determines variability of the model. The corresponding deterministic process on $\mathcal{R}(\kappa)$ with drift $(1-\varrho)\delta$ has value function,

$$J(w;\kappa,\varrho) = \frac{1}{1-\varrho}J(w;\kappa), \quad \kappa \ge 0,\, 0 \le \varrho < 1,\, w \in \mathcal{R}(\kappa).$$

Proposition 5.1: Suppose that (A1)–(A3) hold, and that $\beta_i = 0$ for all i. Then, the TAVC for the process $\widehat{\boldsymbol{W}}$ defined in (65) is non-zero, and satisfies

$$\gamma^2(\varrho,\kappa) = \frac{\kappa^3}{(1-\varrho)^4}\gamma^2(0,1)\,, \quad 0 \le \varrho < 1\,,\ \kappa > 0.$$

PROOF: This follows from the following scaling arguments. Consider first a simple spatial scaling: Let $\widehat{\boldsymbol{W}}$ denote the process with $\varrho = 0$, $\kappa > 0$ arbitrary. For each $\varepsilon > 0$, the process $W^\varepsilon(t) := \varepsilon\widehat{W}(t)$, $t \ge 0$, is a version of (65) with $\varrho = 1-\varepsilon$, and $\kappa' = \varepsilon^2\kappa$. Consequently,

$$\varepsilon^2\gamma^2(0,\kappa) = \gamma^2(1-\varepsilon,\varepsilon^2\kappa), \quad \varepsilon > 0,\ \kappa > 0.$$

Consider next a temporal-spatial scaling: If $\widehat{\boldsymbol{W}}$ is the solution to (65) on $\mathcal{R}$, then $W^\varepsilon(t) := \varepsilon\widehat{W}(t/\varepsilon)$, $t \ge 0$, is also a solution to (65) with ϱ unchanged, and variability parameter $\kappa' = \varepsilon\kappa$. This combined with the definition (63) leads to the identity,

$$\varepsilon^3\gamma^2(\varrho,\kappa) = \gamma^2(\varrho,\varepsilon\kappa)\,, \quad 0 \le \varrho < 1\,,\ \kappa > 0\,,\ \varepsilon > 0.$$

Combining these bounds gives the desired relation.

Finally, positivity of the TAVC follows from the representation (64). □

To investigate the potential for variance reduction we consider the following generalization of the smoothed estimator of [10]:

$$\widehat{\eta}_s(T) = \frac{1}{T}\int_0^T b_{\text{CBM}}(\widehat{W}(t))\, dt\,, \quad T > 0\,, \quad (66)$$

where b_{CBM} is defined in (56) when $\kappa = 1$ and $\varrho = 0$. For arbitrary $0 \le \varrho < 1$, $\kappa > 0$, we set

$$b_{\text{CBM}}(w;\varrho,\kappa) = \Bigl(\frac{\kappa}{1-\varrho}\Bigr) b_{\text{CBM}}(\kappa^{-1}w), \quad w \in \mathcal{R}(\kappa), \quad (67)$$

which is precisely as before, $b_{\text{CBM}} := \bar{c} + \mathcal{D}J$. This definition is entirely analogous to the shadow function used in (18) for the CRW model.

In the next result we demonstrate that the *order* of the TAVC is not improved in the smoothed estimator, unless b_{CBM} is constant on $\mathcal{R}$.

Proposition 5.2: Suppose that (A1)–(A3) hold. Then, the following hold for the process $\widehat{\boldsymbol{W}}$ defined in (65):

(i) The smoothed estimator is consistent. That is,

$$\widehat{\eta}_s(T) \to \widehat{\eta}\,, \qquad T \to \infty.$$

(ii) When $\kappa = 1$ and $\varrho = 0$, the TAVC is given by the mean quadratic-variation of the martingale defined in (60): For each $T > 0$,

$$\gamma_s^2(0,1) = \frac{1}{T}\mathsf{E}_\pi\big[\big(M_{\mathcal{E}}(T)\big)^2\big]. \tag{68}$$

This is zero if and only if b_{CBM} is constant on $\mathcal{R}$, in which case $b_{\text{CBM}}(w) \equiv \widehat{\eta}$.

(iii) If $\beta_i = 0$ for all i, then the TAVC again satisfies,

$$\gamma_s^2(\varrho,\kappa) = \frac{\kappa^3}{(1-\varrho)^4}\gamma_s^2(0,1)\,, \quad 0 \le \varrho < 1\,,\ \kappa > 0.$$

PROOF: Consistency follows from Theorem 4.7 (recall that $\widehat{\eta} := \pi(\bar{c}) = \pi(b_{\text{CBM}})$.) The representation (68) is obtained just as (64) is obtained for the standard estimator.

The scaling result (iii) follows from (67) and the same arguments used in the proof of Proposition 5.1. ☐

While the conclusion of Proposition 5.2 (iii) is negative, in numerical experiments such as illustrated in Figure 5 we have seen dramatic variance reductions.

An alternative way to evaluate a simulation algorithm is through confidence bounds. Given $\varepsilon > 0$ we are interested in bounding the error probability $\mathsf{P}\{|\widehat{\eta}_s(t) - \widehat{\eta}| \ge \varepsilon\}$. Typically one seeks a 'Large Deviation' bound so that this probability decays to zero exponentially fast as $t \to \infty$. However, even for the simple M/M/1 queue it is shown in [11] that the Large Deviations Principle fails.

The story is very different when using the smoothed estimator. The proof of Theorem 5.3 follows from [55, Theorem 6.3 and Theorem 6.5] combined with Theorem 4.7. The 'rate function' Λ^* is the convex dual of the usual log-moment generating function. Without loss of generality we restrict to $\kappa = 1$.

Theorem 5.3: Suppose that the CBM model on $\mathcal{R}$ satisfies (A1)–(A3). If b_{CBM} is not constant on $\mathcal{R}$, then the smoothed estimator (66) satisfies the following exact Large Deviations Principle. There is $\varepsilon_0 > 0$, and functions $H\colon [-\varepsilon_0,\varepsilon_0]\times\mathcal{R} \to \mathbb{R}_+$, $\Lambda^*\colon \mathbb{R} \to \mathbb{R}_+$, such that for each $0 < \varepsilon \le \varepsilon_0$ and initial condition $w \in \mathcal{R}$,

$$\mathsf{P}\{\widehat{\eta}_s(T) \ge \widehat{\eta} + \varepsilon\} \ \sim\ \frac{1}{\sqrt{2\pi T}} H(\varepsilon, w) e^{-\Lambda^*(\varepsilon)T}, \quad T \to \infty.$$

An analogous bound is obtained for $\mathsf{P}\{\widehat{\eta}_s(T) \le \widehat{\eta} - \varepsilon\}$. ☐

VI. LINEAR PROGRAM BOUNDS

We now show that smoothness of the fluid value-function leads to an algorithmic approach to performance approximation, based on the linear-programming approaches of [1], [2], [3], [4], [5], [8].

Suppose that $\bar{c}$ and the region $\mathcal{R}$ satisfy (A1)–(A3). Let $\{R_j : i = 1,\dots,n_1\}$ denote open, connected polyhedral regions satisfying the following: The function b_{CBM} given in (56) is constant on each R_j, $\bar{c}$ is linear on R_j, and $\mathcal{R} =$ closure $(\cup R_j)$.

We consider a family of continuous, piecewise linear functions $\{\bar{k}^i : 1 \le i \le n_2\}$. It is assumed that each of these functions is linear on each of the sets $\{R_j\}$. Consequently, the assumptions of Theorem 4.7 hold: Letting $\{J^i : 1 \le i \le n_2\}$ denote the associated C^1 value functions, and setting $b^i_{\text{CBM}} = \mathcal{D}J^i + \bar{k}^i$, we obtain the identity $\pi(\bar{k}^i) = \pi(b^i_{\text{CBM}})$ for each i. These identities are interpreted as equality constraints below.

The variables in the linear program are defined for $1 \le i \le n$, $1 \le j \le n_1$, by

$$P_j = \pi(R_j), \qquad Y_{ij} = \mathsf{E}_\pi[\widehat{W}_i(t)\mathbb{I}_{R_j}].$$

We have several constraints:

(a) *Mass constraints*: $P_j \ge 0$ for each j, and $\sum P_j = 1$.

(b) *Region constraints*: For example, $Y_{1j} \ge Y_{2j}$ if $\widehat{w}_1 \ge \widehat{w}_2$ within region R_j.

(c) *Value function constraints*: For some constants $\{a_{ij}\} \subset \mathbb{R}$ and vectors $\{\varpi^{ij} : 1 \le i \le n_2\,,\ 1 \le j \le n\} \subset \mathbb{R}^n$ we have the representations for any $1 \le j \le n_1$, $1 \le i \le n_2$,

$$b^i_{\text{CBM}}(w) = a_{ij}; \quad \bar{k}^i(w) = \langle \varpi^{ij}, w\rangle, \qquad w \in R_j.$$

Letting $Y^j = (Y_{1j},\dots,Y_{nj})^{\scriptscriptstyle T} \in \mathbb{R}^n$, $1 \le j \le n_1$, we obtain from Theorem 4.7, for each $i \in \{1,\dots,n_2\}$,

$$\sum_{j=1}^{n_1} \langle \varpi^{ij}, Y^j\rangle = \pi(\bar{k}^i) = \pi(b^i_{\text{CBM}}) = \sum_{j=1}^{n_1} a_{ij} P_j\,. \tag{69}$$

(d) *Objective function*: There is $d \in \mathbb{R}^{n\times n_1}$ such that $\widehat{\eta} := \pi(\bar{c}) = \sum d_{ij} Y_{ij}$.

We illustrate this construction using the CBM workload model for the KSRS network. We assume that the covariance matrix satisfies $\Sigma_{11} = \Sigma_{22} > 0$, and that $\delta = (\delta_1,\delta_1)^{\scriptscriptstyle T} > 0$.

Consider first the affine policy (34) with $\mathcal{R} = \{w \in \mathsf{W} : w_1/3 \le w_2 \le 3w_1\}$. The cost function restricted to $\mathcal{R}$ is the same in Cases I and II, and the common value function shown in (14) is purely quadratic on $\mathcal{R}$. Consequently, in this case we have $h = J$, and

$$\widehat{\eta} = \pi(b_{\text{CBM}}) = b_{\text{CBM}} = \tfrac{1}{8}\delta_1^{-1}(3\Sigma_{11} - \Sigma_{12}). \tag{70}$$

Consider now the minimal process on $\mathsf{W} = \mathbb{R}^2_+$ in Case I. The function b_{CBM} is not constant, so it is not obvious that we can compute $\widehat{\eta}$ exactly using these techniques when $\mathcal{R} = \mathsf{W}$.

To construct an LP we restrict to the following specifications: $n_1 = 3$, with $\{R_i : i = 1,2,3\}$ as shown in Figure 3, and $n_2 = 2$, with $\bar{k}^1(w) = w_1 + w_2$ and $\bar{k}^2(w) = \max(\frac{1}{3}w_1, \frac{1}{3}w_2, \frac{1}{4}(w_1 + w_4))$.

We thus obtain the following constraints:

(a) *Mass constraints*: $P_1 + P_2 + P_3 = 1$

(b) *Region constraints*: We have $Y_{ij} \geq 0$ for all i, j since $\mathsf{W} \subset \mathbb{R}^2_+$. Moreover, on considering the structure of the sets $\{R_i\}$ we obtain, $3Y_{21} \leq Y_{11}$, and $3Y_{13} \leq Y_{23}$. In addition, there are numerous symmetry constraints. For example, $P_1 = P_3$, and $Y_{12} = Y_{22}$ since $\delta_1 = \delta_2$ and $\Sigma_{11} = \Sigma_{22}$.

(c) *Value function constraints*: The value function J^1 is a pure quadratic. In fact, if $k(w) = \langle \bar{k}^1, w \rangle$ is any linear function on W, then

$$J(w) = \tfrac{1}{2}\delta_1^{-1}(\bar{k}_1^1 w_1^2 + \bar{k}_2^1 w_2^2)\,, \quad w \in \mathsf{W}.$$

We conclude from Theorem 4.7 that $\mathsf{E}_\pi[\widehat{W}_1(t)] = \mathsf{E}_\pi[\widehat{W}_2(t)] = \frac{1}{2}\delta_1^{-1}\Sigma_{11}$. The value function J^2 is given in (13). The identity (69) then implies the equality constraint, $\frac{1}{3}Y_{11} + \frac{1}{4}(Y_{12} + Y_{22}) + \frac{1}{3}Y_{23} = \delta_1^{-1}\big[\frac{1}{6}\Sigma_{11}P_1 + \frac{1}{8}(3\Sigma_{11} - \Sigma_{12})P_2 + \frac{1}{6}\Sigma_{22}P_3\big]$.

(d) *Objective function*: In Case I we have,

$$\mathsf{E}_\pi[\bar{c}(\widehat{W}(t))] = \tfrac{1}{3}Y_{11} + \tfrac{1}{4}(Y_{12} + Y_{22}) + \tfrac{1}{3}Y_{23}.$$

We conclude with results from one numerical experiment in Case I, using parameters consistent with the values used in the simulation illustrated in Figure 5 for the controlled random walk model (2). The first-order parameters were scaled as follows:

$$\mu = K[1, 1/3, 1, 1/3], \quad \alpha = \tfrac{1}{4}K\rho_\bullet[1, 0, 1, 0],$$

where K is chosen so that $\sum(\mu_i + \alpha_i) = 1$, and $\rho_\bullet = 0.9$. The effective cost was similarly scaled, $\bar{c}(w) = K\max(w_1/3, w_2/3, (w_1 + w_4)/4)$ for $w \in \mathbb{R}^2_+$.

The random variables $\{S_i(k), A_i(k)\}$ used in (2) were taken mutually independent, with the variance of each random variable equal to its mean. The steady-state covariance is approximated by,

$$\mathsf{Cov}\,[W(k+1) - W(k)] \approx \begin{bmatrix} 13.8042 & 4.2075 \\ 4.2075 & 13.8042 \end{bmatrix},$$

where $W(k) := \Xi Q(k)$. This is exact when $\mathsf{E}[U(k)] = [0.25, 0.75, 0.25, 0.75]^{\scriptscriptstyle T}$.

The linear program constructed for a CBM model is designed to approximate the CRW model: Σ was taken equal to the covariance matrix given above, and $\delta_1 = \delta_2 = 1 - \rho_\bullet = 0.1$. Solving the resulting linear program then gives the following bounds on the steady-state cost, and corresponding occupation probabilities:

$$\begin{aligned} 11.0729 \leq \mathsf{E}_\pi[\bar{c}(\widehat{W}(t))] \leq 14.7638 \\ 0.4895 \leq \mathsf{P}_\pi(\widehat{W}(t) \in R_2) \leq 0.9790 \end{aligned} \tag{71}$$

Although these bounds apply to the CBM model, they roughly approximate the estimated value of $\mathsf{E}_\pi[c(Q(k))] \approx 18$ for the CRW model obtained in the simulation illustrated in Figure 5 [10].

For the same parameters with the region $\mathcal{R} := R_2 = \{w_1/3 \leq w_2 \leq 3w_1\}$ we obtain from (70),

$$\mathsf{E}_\pi[\bar{c}(\widehat{W}(t))] = 14.9218\,.$$

The steady-state mean for the process restricted to the region R_2 is strictly greater than the upper bound (71) obtained for the minimal process on W. This is consistent with the fact that the minimal process on W is optimal whenever $\bar{c}$ is monotone.

VII. Conclusions

The fluid value-function has several desirable properties that provide new algorithms for simulation and performance bounds in stochastic networks, and improved motivation for the application of recently-proposed algorithms.

There are several immediate extensions to be considered.

(i) In [7] the fluid value-function is proposed as an initialization in value iteration to compute optimal policies. It would be worthwhile to revisit this approach for both CBM and CRW network models.

(ii) We have not yet considered examples to test the linear programming techniques on CRW network models.

(iii) In many applications a linear cost function is not appropriate. For example, it may be desirable to obtain bounds on the moment generating function $\mathsf{E}_\pi[\exp(\beta^{\scriptscriptstyle T}\widehat{W}(t))]$ for $\beta \in \mathbb{R}^n$.

(iv) Many of the ideas in this paper may find extension to entirely different areas. In particular, the ODE method commonly used in the analysis of stochastic approximation is completely analogous to the use of a fluid model in the stability analysis of a stochastic network [60]. We are currently investigating the potential for variance reduction in stochastic approximation, and related learning algorithms.

Acknowledgments

This project grew out of on-going discussions with Shane Henderson on simulation. In particular, the C^1 property of the value function was first conjectured by Henderson.

Vivek Borkar also provided much inspiration in discussions on large deviations and on generalizations of Itô's formula. Ruth Williams provided numerous suggestions for clarifying the presentation.

All of this support is gratefully acknowledged.

References

[1] S. Kumar and P. R. Kumar, "Performance bounds for queueing networks and scheduling policies," *IEEE Trans. Automat. Control*, vol. AC-39, pp. 1600–1611, August 1994.

[2] D. Bertsimas, I. Paschalidis, and J. N. Tsitsiklis, "Optimization of multiclass queueing networks: polyhedral and nonlinear characterizations of achievable performance," *Ann. Appl. Probab.*, vol. 4, pp. 43–75, 1994.

[3] P. R. Kumar and S. P. Meyn, "Duality and linear programs for stability and performance analysis queueing networks and scheduling policies," *IEEE Trans. Automat. Control*, vol. 41, no. 1, pp. 4–17, 1996.

[4] E. Schwerer, "A linear programming approach to the steady-state analysis of Markov processes," Ph.D. dissertation, Stanford University, 1997.

[5] J. R. Morrison and P. R. Kumar, "New linear program performance bounds for queueing networks," *J. Optim. Theory Appl.*, vol. 100, no. 3, pp. 575–597, 1999.

[6] S. P. Meyn, "The policy iteration algorithm for average reward Markov decision processes with general state space," *IEEE Trans. Automat. Control*, vol. 42, no. 12, pp. 1663–1680, 1997.

[7] R.-R. Chen and S. P. Meyn, "Value iteration and optimization of multiclass queueing networks," *Queueing Systems: Theory and Applications*, vol. 32, no. 1-3, pp. 65–97, 1999.

[8] M. H. Veatch, "Approximate dynamic programming for networks: Fluid models and constraint reduction," 2004, submitted for publication.

[9] S. G. Henderson and P. W. Glynn, "Approximating martingales for variance reduction in Markov process simulation," *Math. Operations Res.*, vol. 27, no. 2, pp. 253–271, 2002.

[10] S. G. Henderson, S. P. Meyn, and V. B. Tadić, "Performance evaluation and policy selection in multiclass networks," *Discrete Event Dynamic Systems: Theory and Applications*, vol. 13, no. 1-2, pp. 149–189, 2003, special issue on learning, optimization and decision making (invited).

[11] S. P. Meyn, "Large deviation asymptotics and control variates for simulating large functions," 2004, submitted for publication.

[12] J. M. Harrison and R. J. Williams, "Brownian models of open queueing networks with homogeneous customer populations," *Stochastics*, vol. 22, no. 2, pp. 77–115, 1987.

[13] J. M. Harrison and J. A. Van Mieghem, "Dynamic control of Brownian networks: state space collapse and equivalent workload formulations," *Ann. Appl. Probab.*, vol. 7, no. 3, pp. 747–771, 1997.

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

[15] M. Bramson and R. J. Williams, "Two workload properties for Brownian networks," *Queueing Systems: Theory and Applications*, vol. 45, no. 3, pp. 191–221, 2003.

[16] S. P. Meyn, "Sequencing and routing in multiclass queueing networks. Part II: Workload relaxations," *SIAM J. Control Optim.*, vol. 42, no. 1, pp. 178–217, 2003.

[17] M. Chen, C. Pandit, and S. P. Meyn, "In search of sensitivity in network optimization," *Queueing Systems: Theory and Applications*, vol. 44, no. 4, pp. 313–363, 2003.

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

[19] S. P. Meyn, "Dynamic safety-stocks for asymptotic optimality in stochastic networks," 2005, to appear, Queueing Systems.

[20] M. I. Reiman, "Open queueing networks in heavy traffic," *Math. Operations Res.*, vol. 9, pp. 441–458, 1984.

[21] H. Chen and A. Mandelbaum, "Stochastic discrete flow networks: diffusion approximations and bottlenecks," *Ann. Probab.*, vol. 19, no. 4, pp. 1463–1519, 1991.

[22] R. J. Williams, "Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse," *Queueing Systems: Theory and Applications*, vol. 30, no. 1-2, pp. 27–88, 1998.

[23] P. Dupuis and R. J. Williams, "Lyapunov functions for semimartingale reflecting Brownian motions," *Ann. Appl. Probab.*, vol. 22, no. 2, pp. 680–702, 1994.

[24] R. Atar, A. Budhiraja, and P. Dupuis, "On positive recurrence of constrained diffusion processes," *Ann. Probab.*, vol. 29, pp. 979–1000, 2001.

[25] S. P. Meyn and R. L. Tweedie, "Stability of Markovian processes III: Foster-Lyapunov criteria for continuous time processes," *Adv. Appl. Probab.*, vol. 25, pp. 518–548, 1993.

[26] ——, *Markov Chains and Stochastic Stability*. London: Springer-Verlag, 1993, online: `http://black.csl.uiuc.edu/~meyn/pages/book.html`.

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

[28] M. Chen, R. Dubrawski, and S. P. Meyn, "Management of demand-driven production systems," *IEEE Transactions on Automatic Control*, vol. 49, no. 2, pp. 686–698, May 2004.

[29] M. H. Veatch, "Fluid analysis of arrival routing," *IEEE Trans. Automat. Control*, vol. AC-46, pp. 1254–1257, 2001.

[30] J. G. Dai and G. Weiss, "A fluid heuristic for minimizing makespan in job shops," *Operations Res.*, vol. 50, no. 4, pp. 692–707, 2002.

[31] N. Laws, "Dynamic routing in queueing networks," Ph.D. dissertation, Cambridge University, Cambridge, UK, 1990.

[32] S. P. Meyn, "Sequencing and routing in multiclass queueing networks. Part I: Feedback regulation," *SIAM J. Control Optim.*, vol. 40, no. 3, pp. 741–776, 2001.

[33] L. M. Wein, "Dynamic scheduling of a multiclass make-to-stock queue," *Operations Res.*, vol. 40, no. 4, pp. 724–735, 1992.

[34] S. Neumayer, I.-K. Cho, and S. Meyn, "Optimization and the price of anarchy in a dynamic newsboy model," 2005, in preparation.

[35] S. Asmussen, "Queueing simulation in heavy traffic," *Math. Operations Res.*, vol. 17, pp. 84–111, 1992.

[36] W. Whitt, "Planning queueing simulations," *Management Science*, vol. 35, pp. 1341–1366, 1994.

[37] P. Glynn and R. Szechtman, "Some new perspectives on the method of control variates," in *Monte Carlo and Quasi-Monte Carlo Methods 2000: Proceedings of a Conference held at Hong Kong Baptist University, Hong Kong SAR, China*, K. Fang, F. Hickernell, and H. Niederreiter, Eds. Berlin: Springer-Verlag, 2002, pp. 27–49.

[38] S. G. Henderson and B. L. Nelson, Eds., *Simulation*, ser. Handbooks in Operations Research and Management Science, XII. Cambridge, MA: Elsevier, 2005.

[39] M. I. Reiman and R. J. Williams, "A boundary property of semimartingale reflecting Brownian motions," *Probab. Theory Related Fields*, vol. 77, no. 1, pp. 87–97, 1988, (correction in **80** (1989), no. 4, 633).

[40] L. M. Taylor and R. J. Williams, "Existence and uniqueness of semimartingale reflecting Brownian motions in an orthant," *Probab. Theory Related Fields*, vol. 96, no. 3, pp. 283–317, 1993.

[41] P. Dupuis and K. Ramanan, "Convex duality and the Skorokhod problem. I, II," *Probab. Theory Related Fields*, vol. 115, no. 2, pp. 153–195, 197–236, 1999.

[42] P. Dupuis and I. Hitoshi, "On Lipschitz continuity of the solution mapping to the Skorokhod problem, with applications," *Stochastics Stochastics Rep.*, vol. 35, no. 1, pp. 31–62, 1991.

[43] S. N. Ethier and T. G. Kurtz, *Markov Processes : Characterization and Convergence*. New York: John Wiley & Sons, 1986.

[44] N. V. Krylov, *Controlled diffusion processes*, ser. Applications of Mathematics. New York: Springer-Verlag, 1980, vol. 14, translated from the Russian by A. B. Aries.

[45] ——, "On a proof of Itô's formula," *Trudy Mat. Inst. Steklov.*, vol. 202, pp. 170–174, 1993.

[46] R. Chitashvili and M. Mania, "Generalized Itô formula and derivation of Bellman's equation," in *Stochastic processes and related topics (Siegmundsberg, 1994)*, ser. Stochastics Monogr. Yverdon: Gordon and Breach, 1996, vol. 10, pp. 1–21.

[47] H. Föllmer and P. Protter, "On Itô's formula for multidimensional Brownian motion," *Probab. Theory Related Fields*, vol. 116, no. 1, pp. 1–20, 2000.

[48] P. W. Glynn and S. P. Meyn, "A Liapounov bound for solutions of the Poisson equation," *Ann. Probab.*, vol. 24, no. 2, pp. 916–931, 1996.

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

[50] D. Q. Mayne, "Control of constrained dynamic systems," *European J. of Control*, vol. 7, pp. 87–99, 2001.

[51] F. Borrelli, *Constrained optimal control of linear and hybrid systems*, ser. Lecture Notes in Control and Information Sciences. Berlin: Springer-Verlag, 2003, vol. 290.

[52] H. Thorisson, *Coupling, stationarity, and regeneration*, ser. Probability and its Applications (New York). New York: Springer-Verlag, 2000.

[53] L. C. G. Rogers and D. Williams, *Diffusions, Markov processes, and martingales. Vol. 2*, ser. Cambridge Mathematical Library. Cambridge: Cambridge University Press, 2000, ito calculus, Reprint of the second (1994) edition.

[54] D. Down, S. P. Meyn, and R. L. Tweedie, "Exponential and uniform ergodicity of Markov processes," *Ann. Probab.*, vol. 23, no. 4, pp. 1671–1691, 1995.

[55] I. Kontoyiannis and S. P. Meyn, "Spectral theory and limit theorems for geometrically ergodic Markov processes," *Ann. Appl. Probab.*, vol. 13, pp. 304–362, 2003, presented at the INFORMS Applied Probability Conference, NYC, July, 2001.

[56] S. P. Meyn and R. L. Tweedie, "Generalized resolvents and Harris recurrence of Markov processes," *Contemporary Mathematics*, vol. 149, pp. 227–250, 1993.

[57] J. M. Harrison and M. I. Reiman, "On the distribution of multidimensional reflected Brownian motion," *SIAM J. Appl. Math.*, vol. 41, no. 2, pp. 345–361, 1981.

[58] J. M. Harrison and R. J. Williams, "Brownian models of feedforward queueing networks: quasireversibility and product form solutions," *Ann. Appl. Probab.*, vol. 2, no. 2, pp. 263–293, 1992.

[59] J. G. Dai and J. M. Harrison, "Reflected Brownian motion in an orthant: Numerical methods for steady-state analysis," *Ann. Appl. Prob.*, vol. 2, pp. 65–86, 1992.

[60] V. S. Borkar and S. P. Meyn, "The O.D.E. method for convergence of stochastic approximation and reinforcement learning," *SIAM J. Control Optim.*, vol. 38, no. 2, pp. 447–469, 2000, (also presented at the *IEEE CDC*, December, 1998).