

SEQUENCING AND ROUTING IN MULTICLASS QUEUEING NETWORKS PART II: WORKLOAD RELAXATIONS*

SEAN P. MEYN[†]

Abstract. Part II continues the development of policy synthesis techniques for multiclass queueing networks based upon a linear fluid model. The following are shown:

(i) A relaxation of the fluid model based on workload leads to an optimization problem of lower dimension. An analogous workload-relaxation is introduced for the stochastic model. These relaxed control problems admit pointwise optimal solutions in many instances.

(ii) A translation to the original fluid model is *almost optimal*, with vanishing relative error as the network load ρ approaches one. It is pointwise optimal after a short transient period, provided a pointwise optimal solution exists for the relaxed control problem.

(iii) A translation of the optimal policy for the fluid model provides a policy for the stochastic network model that is almost optimal in heavy traffic, over all solutions to the relaxed stochastic model, again with vanishing relative error. The regret is of order $|\log(1 - \rho)|$.

Key words. queueing networks, routing, scheduling, optimal control

AMS subject classifications. Primary 90B35, 68M20, 90B15; Secondary 93E20, 60J20

PII. S036301290138376X

1. Introduction. Variability in a queueing system may have significant impact on performance. Kingman's bound implies that in heavy traffic, when the load ρ is close to unity, even small variations in service times or interarrival times can lead to long delays [3]. This is one reason for the much-publicized difficulties in the air-traffic, highway, and power industries, where real-life networks in heavy traffic are experienced each day by pilots, passengers, and home-owners (see, e.g., [45, 5]).

Variability often plays a smaller role in *relative performance* in network models, when comparing two candidate policies for network regulation (although this depends upon the specific network topology and the performance metric under consideration). For this reason, variability often plays a minor role in many aspects of network design and analysis. Stability of a network under a particular policy is determined by first order statistics (mean arrival and service rates and routing probabilities), except in pathological examples [8, 10, 9, 34, 36]. Part I primarily concerns policy synthesis in queueing networks. It is shown that robustly stabilizing policies can be constructed by appropriately translating a policy for the associated linear fluid model, which is defined only by first order statistics (see [36] for a bibliography).

This paper continues the development of part I. We focus primarily on policy synthesis and network optimization because of the intrinsic importance of these issues, and because it is likely that a deeper understanding will lead to improved methods for addressing many other issues in design, such as performance approximation and network planning. Among the issues not addressed in [36, 35] are the following:

(i) *The role of variability in design.* It is known that an understanding of variability is important in the determination of safety stocks to prevent unwanted idleness. Is this the only use of high order statistical information in policy synthesis?

*Received by the editors January 11, 2001; accepted for publication (in revised form) September 7, 2002; published electronically March 26, 2003. This work was supported in part by NSF grants DMI 00 85165 and ECS 99 72957.

<http://www.siam.org/journals/sicon/42-1/38376.html>

[†]Coordinated Science Laboratory and the University of Illinois, 1308 W. Main Street, Urbana, IL 61801 (s-meyn@uiuc.edu, <http://black1.csl.uiuc.edu:80/~meyn>).

(ii) *Complexity management.* For example, is it possible to construct optimal policies for the fluid model when the network is large?

(iii) *Performance validation.* Will a translation lead to an optimal allocation for the physical network?

Some answers to these questions are provided here.

A series of recent papers in the stochastic network literature show that a combination of “resource pooling” and “state space collapse” occur in heavy traffic, where $\rho \sim 1$ [38, 39, 43, 29, 25, 4]. See also the recent monographs [6, 30]. State space collapse can transform a network with hundreds of buffers into a far simpler model that retains most of the essential information required for the design of efficient policies. All of these prior results are based on a reflected Brownian motion (RBM) model to approximate the network of interest. This approach is not pursued here for several reasons: Technicalities arising in a proof of weak convergence to an RBM model are avoided, and as pointed out in part I, it is not necessary to assume that the network is *balanced* (i.e., loads at all stations are comparable). This allows significantly greater flexibility in modelling. In this paper we also find that the “Brownian motion scaling” may wash away too many details. By avoiding any scaling, relative bounds on performance are obtained that are far stronger than reported previously in any examples.

As in part I, the primary model considered here is the linear fluid model (2.5). One of the main contributions of the present paper is to introduce a *workload-relaxation* of the fluid model that may be viewed as a generalization of state space collapse, as formulated in the aforementioned references. The significant model reduction obtained in a workload-relaxation provides a framework for addressing many aspects of (i)–(iii).

We show in particular that very strong solidarity exists between respective optimal control solutions. Let c denote a norm on the state space of buffer-levels $\mathbf{X} := \mathbb{R}_+^\ell$ —in the results below we eventually specialize to piecewise linear functions on \mathbf{X} . Suppose that \mathbf{Q} is any queue length process evolving on \mathbf{X} defined by some admissible policy. Kingman’s bound will then give a steady-state bound of the form

$$\mathbb{E}[c(\mathbf{Q}(t))] \geq O\left(\frac{1}{1-\rho}\right).$$

Suppose that \mathbf{Q}° is the process on \mathbf{X} obtained through tracking the optimal fluid model trajectories, as described below. Under general geometric conditions (including uniqueness of solutions to the fluid-model optimal-control problem), we show in Theorem 4.3 that \mathbf{Q}° is *approximately optimal, with logarithmic regret*: as $\rho \uparrow 1$,

$$\frac{1}{T} \int_0^T c(\mathbf{Q}^\circ(t; x)) dt \leq \frac{1}{T} \int_0^T c(\mathbf{Q}(t; x)) dt + O(\log((1-\rho)^{-1})), \quad 0 \leq T \leq \frac{1}{(1-\rho)^3},$$

where \mathbf{Q} is any other solution. We also find that no formulation of sample-path optimality is feasible in heavy traffic under complementary geometric conditions. Consequently, extensions of the results reported here require comparison of a mean-performance metric, rather than sample path bounds (see [7] for recent results in this direction).

The remainder of the paper is organized as follows. Section 2 provides a description of a stochastic network model and the linear fluid model. A reduced order model based on “workload-relaxation” is developed in section 3, and optimal policies for the relaxation are constructed.

Section 4 concerns models in heavy traffic, where $\rho \sim 1$. A policy is constructed based on a translation of the optimal solution to the relaxed fluid-model optimal-control problem. It is shown that this translation is almost optimal for the original fluid model, with bounded error as the system load approaches unity. When a reflected Brownian motion limit exists in heavy traffic, then this “state space collapse” coincides with that observed in the aforementioned references. Similar results hold for a general stochastic model: it is shown that this policy is approximately optimal for a stochastic model, with logarithmic regret, over all solutions to a relaxation of the associated stochastic optimal-control problem.

Section 5 contains conclusions and poses various possible extensions.

2. Models and control. As in [36], this paper is based on a stochastic, bursty model, and a linear fluid model that may be interpreted as a scaled version of its bursty counterpart.

2.1. The stochastic model. The network model described here is a version of the stochastic processing network developed in [23, 24]. We denote by \mathbf{Q} the stochastic process evolving on $\mathbf{X} = \mathbb{R}_+^\ell$ whose components indicate buffer levels for the stochastic network model. For example, the network shown in Figure 1 is a simple manufacturing model in which $\ell = 16$, and four of these buffers are *virtual*, corresponding to backlog or excess inventory.

For a given initial condition $Q(0; x) = x \in \mathbf{X}$ the dynamics of \mathbf{Q} are expressed

$$(2.1) \quad Q(t; x) = x - S(Z(t; x)) + R(Z(t; x)) + A(t), \quad t \geq 0.$$

The vector-valued stochastic process \mathbf{Z} is the *allocation* (or *control*) evolving on $\mathbb{R}_+^{\ell_u}$ for some integer ℓ_u . The i th component $Z_i(t; x)$ gives the cumulative time that the activity i has run up to time t , $1 \leq i \leq \ell_u$. Activities may include a combination of *sequencing* of various jobs at a particular station and *routing* those jobs to other stations once service is completed. Several examples are given in [36].

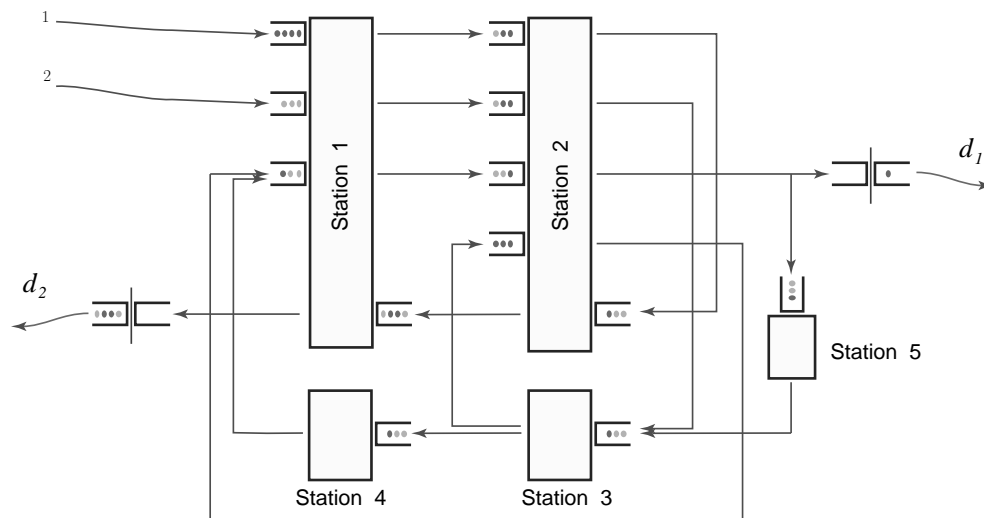


FIG. 1. A network with many buffers, controlled routing, uncontrolled routing, multiple demands, and virtual buffers.

The allocation rates are subject to linear constraints

$$(2.2) \quad Z(t; x) - Z(s; x) \geq \boldsymbol{\theta}, \quad C[Z(t; x) - Z(s; x)] \leq [t - s]\mathbf{1}, \quad t \geq s \geq 0,$$

where the *constituency matrix* C is an $\ell_m \times \ell_u$ matrix with binary entries; $\boldsymbol{\theta}$ is a vector of zeros; and $\mathbf{1}$ is a vector of ones. The *i th resource* \mathcal{R}_i is defined to be the set of activities j such that $C_{ij} = 1$. The constraint (2.2) expresses the condition that resources are shared, and they are limited.

The process \mathbf{A} may denote a combination of exogenous arrivals to the network and exogenous demands for materials *from* the network. The function $S(\cdot)$ represents possibly random service times, and the function $R(\cdot)$ represents the effects of a combination of possibly uncontrolled, possibly random routing and random service times.

Specific statistical assumptions on $\{\mathbf{A}, \mathbf{R}, \mathbf{S}\}$ are given in section 4.2 where the stochastic model is considered in detail. Many of the variables $\{A_i(\cdot), R_i(\cdot), S_i(\cdot)\}$ will be null in general, and they are typically highly correlated.

The average-cost optimization problem is concerned with minimizing the long-run average cost,

$$(2.3) \quad \Gamma(x) = \limsup_{T \rightarrow \infty} \mathbb{E} \left[\frac{1}{T} \int_0^T c(Q(t; x)) dt \right],$$

subject to the constraints given above, where $c: \mathbb{R}^\ell \rightarrow \mathbb{R}_+$ is a convex function that vanishes only at the origin. In section 4.2 we consider generalizations in which $c(\cdot)$ is also a function of \mathbf{Z} . In this case the cost function may be chosen to reflect the desire to maximize utilization of some resources, while minimizing utilization of others.

It is clear that an exact optimal solution to (2.3) will not be found except in very special cases.

2.2. The linear fluid model. Assumption S, to be imposed in section 4.2, implies that the law of large numbers holds: For some $\ell \times \ell_u$ matrix B , a vector $\alpha \in \mathbb{R}_+^\ell$, and any $z \in \mathbb{R}_+^{\ell_u}$,

$$(2.4) \quad Bz = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T [-S(zt) + R(zt)] dt, \quad \alpha = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T A(t) dt.$$

This provided motivation in [36] for the fluid analogue of (2.1) given by

$$q(t; x) = x + Bz(t; x) + \alpha t, \quad t \geq 0.$$

The vector $\zeta(t; x) = \frac{d}{dt} z(t; x)$ denotes allocation rates, and $q(t; x)$ is a vector of buffer levels. This is also expressed as the controlled, linear ordinary differential equation

$$(2.5) \quad \frac{d}{dt} q(t; x) = B\zeta(t; x) + \alpha, \quad t \geq 0, \quad q(0; x) = x,$$

where throughout the paper the symbol “ $\frac{d}{dt}$ ” denotes a *right derivative*.

It is convenient to envision (2.5) as a differential inclusion:

- (i) The state \mathbf{q} is constrained to evolve in the polyhedron $\mathbf{X} = \mathbb{R}_+^\ell$.
- (ii) The allocation rates $\boldsymbol{\zeta}$ evolve in a polyhedron $\mathbf{U} \subseteq \mathbb{R}_+^{\ell_u}$, defined by

$$\mathbf{U} := \{\boldsymbol{\zeta} \in \mathbb{R}^{\ell_u} : \boldsymbol{\zeta} \geq \boldsymbol{\theta}, C\boldsymbol{\zeta} \leq \mathbf{1}\}.$$

(iii) The velocity $\frac{d}{dt}q$ is constrained to lie in the polyhedron

$$\mathbf{V} := \{v = B\zeta + \alpha : \zeta \in \mathbf{U}\}.$$

The assumptions below imply that the network can be controlled so that, starting empty, it will remain empty. This means that \mathbf{V} contains the origin, or equivalently, there exists at least one solution $\zeta^{ss} \in \mathbf{U}$ to the equilibrium equation

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

Section 2.3 is concerned with the existence of equilibria and simple formulations of *optimality* for ζ^{ss} .

Two dynamic optimization problems are singled out because of their mathematical and economic interest:

Time-optimal control. For any initial condition $q(0) = x$, find an allocation \mathbf{z} that minimizes

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

The minimal draining time is denoted $T^*(x)$, with the convention that the minimum over an empty set is interpreted as infinity.

Total-cost optimal control. For any initial condition $q(0) = x$, find an allocation \mathbf{z} that minimizes

$$(2.6) \quad J(x) = \int_0^T c(q(t; x)) dt.$$

We consider primarily the infinite-horizon case in which $T = \infty$, and in this case we let J^* denote the “optimal cost” (i.e., the infimum over all policies).

The fluid model is called *stabilizable* if $T^*(x) < \infty$ for all $x \in \mathbf{X}$. If the model is stabilizable, then there exists a time-optimal allocation that is *linear*. For any $x \in \mathbf{X}$, if \mathbf{z} is any time-optimal allocation, then we write

$$(2.7) \quad \zeta^\circ = \frac{z(T^*(x); x)}{T^*(x)} \in \mathbf{U}.$$

The allocation $z^\circ(t; x) = t\zeta^\circ$, $0 \leq t \leq T^*(x)$, is evidently feasible and time-optimal. This linear policy and stochastic translations are considered in [11], and generalizations are treated in [17, 14].

The infinite-horizon cost criterion is more closely aligned with the average-cost optimization problem. Computing J^* and an optimal allocation z^* can be formulated as an infinite-dimensional linear program when the cost c is piecewise linear [37]. Algorithms are available that solve this control problem for models of moderate complexity [32, 42].

In the remainder of the paper we take c to be piecewise linear, of the form

$$(2.8) \quad c(x) = \max_{1 \leq i \leq \ell_c} \langle c^i, x \rangle, \quad x \in \mathbb{R}^\ell,$$

where $\langle \cdot, \cdot \rangle$ denotes the usual inner product on \mathbb{R}^ℓ . We can approximate any norm on \mathbb{R}^ℓ through an appropriate choice of $\{c^i\} \subset \mathbb{R}^\ell$. A lower bound on performance,

at a specific time t , given a specific initial condition $x \in \mathbf{X}$, is found by solving the linear program

$$\begin{aligned}
 (2.9) \quad & \min \quad \gamma \\
 & \text{subject to} \quad \gamma \geq \langle c^i, y \rangle, \quad 1 \leq i \leq \ell_c, \\
 & \quad \quad \quad y = x + Bz + \alpha t, \\
 & \quad \quad \quad Cz \leq t\mathbf{1}, \\
 & \quad \quad \quad y, z \geq \boldsymbol{\theta}.
 \end{aligned}$$

We denote the value of this linear program by $\underline{c}^*(t; x)$. A feasible state trajectory q^* starting from x is called *pointwise optimal* if $c(q^*(t; x)) = \underline{c}^*(t; x)$ for every t . A pointwise optimal trajectory is always time-optimal, and it is also *greedy*: The derivative $\frac{d}{dt}c(q(t; x))$ is minimized over all allocation rates at each time t .

It is rare to find a model for which a pointwise optimal solution exists from each initial condition. However, in section 3 general conditions are formulated which ensure that $c(q^*(t; x)) = \underline{c}^*(t; x)$ for all t following a short transient period.

A first step towards optimization is stabilizability: When are T^* and J^* finite-valued? What is the network load?

2.3. Capacity and time-optimal control. If the fluid model is stabilizable, then the origin is an equilibrium for the model, which means that $\boldsymbol{\theta} \in \mathbf{V}$. We let \mathbf{U}_{ss} denote the set of all allocation rates that achieve this: $\mathbf{U}_{\text{ss}} = \{\zeta : B\zeta + \alpha = \boldsymbol{\theta}\}$. In the classical *scheduling* problem there is a unique activity associated with each buffer. This implies that the matrix B is square, and stabilizability ensures that B is full-rank. It then follows that \mathbf{U}_{ss} contains a unique vector of steady-state allocation rates given by $\zeta^{\text{ss}} := -B^{-1}\alpha$. We then define the *vector load* by

$$(2.10) \quad \boldsymbol{\rho} = (\rho_1, \dots, \rho_\ell)^T = -CB^{-1}\alpha = C\zeta^{\text{ss}},$$

and the *system load* is the maximum, $\rho = \max_i \rho_i$.

In other models the matrix B may not be square. The set of equilibrium rates \mathbf{U}_{ss} may be large, and some may impose a greater “load” on the system than others. The following is taken from [23], following [29, 22]. The network load ρ is defined as the solution to the linear program

$$\begin{aligned}
 (2.11) \quad & \min \quad \rho \\
 & \text{subject to} \quad B\zeta + \alpha = \boldsymbol{\theta}, \\
 & \quad \quad \quad C\zeta \leq \rho\mathbf{1}, \\
 & \quad \quad \quad \zeta \geq \boldsymbol{\theta}.
 \end{aligned}$$

The idea is that we consider all allocation rates ζ^{ss} that provide an equilibrium and choose among these the one that has minimal overall impact on the system in the sense that $\max_i [C\zeta^{\text{ss}}]_i$ is smallest.

Closely related is the linear program defining the minimal draining time

$$\begin{aligned}
 (2.12) \quad & \min \quad T \\
 & \text{subject to} \quad x + Bz + \alpha T = \boldsymbol{\theta}, \\
 & \quad \quad \quad Cz \leq T\mathbf{1}, \\
 & \quad \quad \quad z \geq \boldsymbol{\theta},
 \end{aligned}$$

where $x \in \mathbb{R}^\ell$ is given. The value of this linear program is equal to $T^*(x)$.

We let $W^*(x)$ denote the minimum time to drain the fluid model for an arrival-free model where $\alpha = \theta$. The definition of load is thus motivated by considering the fluid model (2.5) without arrivals: on comparing (2.12) and (2.11) it is seen that $\rho = W^*(\alpha)$. Thus, if α units of material arrives at the network in one second, the *system load* is the amount of time required to clear this material, given that no other material arrives.

Alternative representations for the minimal emptying times are found through a representation of the velocity set \mathbf{V} . Let \mathbf{V}_0 denote the velocity set for the arrival-free model:

$$(2.13) \quad \mathbf{V}_0 := \{v - \alpha : v \in \mathbf{V}\} = \{B\zeta : \zeta \in \mathbf{U}\}.$$

THEOREM 2.1. *The sets \mathbf{V}_0, \mathbf{V} are described by the intersection of half-spaces: There exists a set of vectors $\{\xi^i : 1 \leq i \leq \ell_v\} \subset \mathbb{R}^{\ell_v}$, for some minimal integer $\ell_v \geq 1$, and binary values $\{b_i : 1 \leq i \leq \ell_v\} \subset \{0, 1\}$ such that the following hold:*

$$(2.14) \quad \mathbf{V}_0 = \{v : \langle \xi^i, v \rangle \geq -b_i, 1 \leq i \leq \ell_v\},$$

$$(2.15) \quad \mathbf{V} = \{v : \langle \xi^i, v \rangle \geq -(b_i - \rho_i), 1 \leq i \leq \ell_v\},$$

where in (2.15) we set $\rho_i = \langle \xi^i, \alpha \rangle$ for $1 \leq i \leq \ell_v$.

Proof. The representation \mathbf{V}_0 in (2.14) follows from the fact that it is a polyhedral subset of \mathbb{R}^{ℓ} containing the origin. The representation for \mathbf{V} then follows from the formula $\mathbf{V} = \{v + \alpha : v \in \mathbf{V}_0\}$ and the definition of $\{\rho_i\}$. \square

The vector ξ^i is called a *workload vector* if $b_i \neq 0$. We denote by ℓ_r the number of distinct workload vectors.

For a given $\alpha \in \mathbb{R}_+^{\ell}$ we assume that the vectors $\{\xi^i\}$ are ordered so that $\rho_1 \geq \dots \geq \rho_{\ell_v}$. Provided the linear program defining ρ is feasible, we see from Theorem 2.2(ii) that, under this ordering, the set of workload vectors is given by $\{\xi^i : 1 \leq i \leq \ell_r\}$ and that the system load defined in (2.11) is equal to ρ_1 .

THEOREM 2.2. *The following hold for the model (2.5), for any given $\alpha \in \mathbb{R}_+^{\ell}$, $x \in \mathbf{X}$:*

(i) *If $\langle \xi^i, x \rangle > 0$ for some $i > \ell_r$, then $W^*(x) = \infty$. Otherwise, the minimal emptying time for the arrival-free model is finite and given by*

$$W^*(x) = \max_{1 \leq i \leq \ell_r} \langle \xi^i, x \rangle.$$

(ii) *Suppose that the constraint set in the linear program (2.11) is nonempty. Then, $\rho_i \leq 0$ for $i > \ell_r$, and the system load can be expressed as*

$$\rho = W^*(\alpha) = \max_{1 \leq i \leq \ell_r} \rho_i.$$

(iii) *If $\langle \xi^i, x \rangle > 0$ and $\rho_i \geq 0$ for some $i > \ell_r$, then $T^*(x) = \infty$. Otherwise, provided $\rho < 1$, the minimal emptying time T^* is finite and given by*

$$T^*(x) = \max_{1 \leq i \leq \ell_v} \frac{\langle \xi^i, x \rangle}{b_i - \rho_i}.$$

(iv) *The model is stabilizable if $\rho < 1$, and $\rho_i < 0$ for $i > \ell_r$. The second condition is automatic if the arrival-free model is stabilizable.*

Proof. Part (i) follows from Theorem 2.1: for $x \neq \theta$, provided $W^*(x) < \infty$,

$$\begin{aligned} W^*(x) &= \min(T > 0 : -T^{-1}x \in V_0) \\ &= \min(T > 0 : \langle \xi^i, -T^{-1}x \rangle \geq b_i, 1 \leq i \leq \ell_v) \\ &= \min(T > 0 : \langle \xi^i, x \rangle \leq b_i T, 1 \leq i \leq \ell_v). \end{aligned}$$

Recall that $b_i = 0$ for $i > \ell_r$. If for some such i we have $\langle \xi^i, x \rangle > 0$, then we see that the constraint set in the minimization is infeasible, and we conclude that $W^*(x)$ cannot be finite. Conversely, if $\langle \xi^i, x \rangle \leq 0$ for $i > \ell_r$, then the equation above gives the desired representation for W^* . This establishes (i), and (iii) follows similarly using the definition $\rho_i := \langle \xi^i, \alpha \rangle$.

The proof of (ii) follows from (i) and the representation $\rho = W^*(\alpha)$, and result (iv) follows directly from (iii). \square

The workload vectors may be interpreted as Lagrange multipliers since they define sensitivity of the optimal draining time with respect to the initial condition x . The following results provide further interpretations. For a given $x \in \mathbb{R}^\ell$, consider the dual of the linear program (2.12)

$$(2.16) \quad \begin{aligned} &\max \quad \langle \xi, x \rangle \\ &\text{subject to} \quad \begin{aligned} B^T \xi + C^T \eta &\geq \theta, \\ -\alpha^T \xi + \mathbf{1}^T \eta &\leq 1, \\ \eta &\geq \theta. \end{aligned} \end{aligned}$$

On considering the extreme points of (2.16), we may express the value of this linear program as a piecewise linear function on the domain $\{x \in \mathbb{R}^\ell : T^*(x) < \infty\}$. Applying Theorem 2.2, we see that these correspond to the vectors $\{\xi^i : 1 \leq i \leq \ell_r\}$ used in the representation of the sets V and V_0 .

In view of this solidarity we denote by $\{(\xi^i, \eta^i) : 1 \leq i \leq \ell_r\}$ the nonzero extreme points of the constraint set in (2.16) when $\alpha = \theta$. For each i we have $\xi^i \in \mathbb{R}^\ell, \eta^i \in \mathbb{R}_+^{\ell_m}$. An interpretation of the vectors $\{\eta^i\}$ is provided in the following proposition.

PROPOSITION 2.3. *Consider the linear program (2.16) with $\alpha = \theta$. If (ξ, η) is an extreme point in the constraint set satisfying $\xi \neq \theta$, then $\eta \in \mathbb{R}_+^{\ell_m}$ satisfies $\langle \mathbf{1}, \eta \rangle = 1$. Consequently, for any $1 \leq i \leq \ell_r$ we may interpret the vector η^i as a probability distribution on the resources $\{1, \dots, \ell_m\}$.*

Proof. Suppose that (ξ, η) is any feasible pair with $0 \leq \langle \mathbf{1}, \eta \rangle < 1$, and $\xi \neq \theta$. Then $(\gamma\xi, \gamma\eta)$ is also feasible for any $0 < \gamma < \langle \mathbf{1}, \eta \rangle^{-1}$, which implies that (ξ, η) cannot be an extreme point. \square

The *workload process* is defined on a fluid scale by

$$(2.17) \quad w(t; x) = \Xi q(t; x), \quad t \geq 0, \quad x \in X,$$

where Ξ denotes the $\ell_r \times \ell$ matrix whose i th row is given by ξ^{iT} .

PROPOSITION 2.4. *The following lower bounds hold:*

- (i) $\langle \xi^i, B\zeta \rangle \geq -b_i, \quad \zeta \in U, \quad 1 \leq i \leq \ell_v;$
- (ii) $\frac{d}{dt} w_i(t; x) \geq -(1 - \rho_i), \quad t \geq 0, \quad 1 \leq i \leq \ell_r.$

Proof. For (i), note that $v_0 := B\zeta$ is a generic element of V_0 , so the result follows from the representation of V_0 in Theorem 2.1. As for (ii), observe that

$$\frac{d}{dt} w_i = \langle \xi^i, v \rangle,$$

where $v := B\zeta + \alpha$ is a generic element of \mathbf{V} . This and Theorem 2.1 again imply the result since $b_i = 1$ for $1 \leq i \leq \ell_r$. \square

We define the i th set of *pooled-resources* by

$$\mathcal{R}_i^\circ := \{j \leq \ell_m : \eta_j^i > 0\}, \quad 1 \leq i \leq \ell_r.$$

Resource j is called a *bottleneck* if $j \in \mathcal{R}_i^\circ$ for some $i \leq \ell_r$, and $\rho_i = \rho$. The following result provides motivation for this terminology.

PROPOSITION 2.5. *For any $1 \leq i \leq \ell_r$, and any $\zeta \in \mathbf{U}$, the following are equivalent:*

- (i) $\langle \xi^i, B\zeta \rangle = -1$,
- (ii) $(C\zeta)_j = 1$ for all $j \in \mathcal{R}_i^\circ$, and ζ satisfies the complementary slackness condition

$$\zeta_j > 0 \implies [B^T \xi^i + C^T \eta^i]_j = 0.$$

Proof. Suppose that (i) holds. Then we may multiply ζ^T times the constraint $B^T \xi + C^T \eta \geq \theta$ in (2.16) to obtain the bound

$$-1 + \langle \eta^i, C\zeta \rangle = \langle \xi^i, B\zeta \rangle + \langle \eta^i, C\zeta \rangle = \langle \zeta, [B^T \xi^i + C^T \eta^i] \rangle \geq 0,$$

and it follows that $\langle \eta^i, C\zeta \rangle \geq 1$. Since the reverse inequality also holds when $\zeta \in \mathbf{U}$, we must have equality:

$$(2.18) \quad -1 + \langle \eta^i, C\zeta \rangle = \langle \zeta, [B^T \xi^i + C^T \eta^i] \rangle = 0.$$

In fact, since η^i is a probability distribution on $\{1, \dots, \ell_u\}$ and $C\zeta \leq \mathbf{1}$, the equality (2.18) implies that $(C\zeta)_j = 1$ for all $j \in \mathcal{R}_i^\circ$. The equation (2.18) also implies the complementary slackness condition in (ii) since $[B^T \xi^i + C^T \eta^i] \geq \theta$, and $\zeta \in \mathbb{R}_+^{\ell_m}$.

Conversely, if (ii) holds, then the complementary slackness condition implies the identity, $\langle \xi^i, B\zeta \rangle + \langle \eta^i, C\zeta \rangle = \langle \zeta, [B^T \xi^i + C^T \eta^i] \rangle = 0$. This combined with the assumption in (ii) that $(C\zeta)_j = 1$ whenever $j \in \mathcal{R}_i^\circ$ (equivalently $\langle \eta^i, C\zeta \rangle = 1$) gives (i) immediately. \square

The workload vectors allow us to define “hot spots” in the network and give some intuition about the structure of good policies. Suppose that at time t the state takes the value $q(t; x) = y$. The i th pooled-resource is a *dynamic bottleneck* at time t if

$$T^*(y) = \langle \xi^i, y \rangle / (1 - \rho_i).$$

An ordinary resource j is called a *dynamic bottleneck* at time t if $j \in \mathcal{R}_i^\circ$ for some $1 \leq i \leq \ell_r$, and pooled-resource i is a dynamic bottleneck. We say that the i th pooled-resource is *working at capacity* at time t if $\langle \xi^i, B\zeta(t) \rangle = -1$.

The following is then immediate from Proposition 2.4, Proposition 2.5, and Theorem 2.2.

THEOREM 2.6. *Suppose that \mathbf{q} is any solution to the fluid-model equations (2.5) starting at $x \in \mathbf{X}$.*

(i) *If \mathbf{q} is time-optimal (so that $q(t; x) = \theta$ for $t \geq T^*(x)$), then each dynamic-bottleneck pooled-resource is working at capacity for each $t < T^*(x)$.*

(ii) *If each dynamic-bottleneck pooled-resource works at capacity for $t < T^*(x)$, then the state trajectory \mathbf{q} is time-optimal.* \square

3. The relaxed control problem. We introduce here a relaxation of the optimal control problem (2.6). The main idea is that, for the purposes of control, only a few of the workload vectors impose serious constraints. A much simpler optimal control problem is obtained by relaxing those constraints corresponding to relatively small load.

3.1. Almost-equivalent workload formulation. For arbitrary $1 \leq n \leq \ell_r$, the n th relaxation of (2.5) is defined as follows. As before, the state space \mathbf{X} is taken as \mathbb{R}_+^ℓ , but the velocity set is given by

$$\widehat{\mathbf{V}} = \{v : \langle \xi^i, v \rangle \geq -(1 - \rho_i), 1 \leq i \leq n\}.$$

An application of Theorem 2.1 establishes the inclusion $\mathbf{V} \subset \widehat{\mathbf{V}}$. It is assumed throughout that $\{\xi^i : 1 \leq i \leq n\}$ are linearly independent vectors.

We denote by \widehat{q} any feasible state trajectory:

$$(3.1) \quad \widehat{q}(0; x) = x, \quad \widehat{q}(t; x) \in \mathbf{X}, \quad \text{and} \quad \frac{q(t; x) - q(s; x)}{t - s} \in \widehat{\mathbf{V}}, \quad 0 \leq s < t.$$

The n th relaxation may also be described in a form analogous to (2.5):

$$(3.2) \quad \frac{d}{dt} \widehat{q}(t; x) = B \widehat{\zeta}(t; x) + \alpha, \quad t \geq 0, \quad \widehat{q}(0; x) = x.$$

The allocation rates in (3.2) are subject to the constraints

$$\widehat{\zeta}(t; x) \in \widehat{\mathbf{U}} := \{\zeta \in \mathbb{R}^{\ell_u} : \widehat{C} \zeta \leq \mathbf{1}\},$$

where $\widehat{C} := -\widehat{\Xi}B$, and $\widehat{\Xi}$ denotes the $n \times \ell$ matrix

$$(3.3) \quad \widehat{\Xi} = [\xi^1 \mid \dots \mid \xi^n]^T.$$

The equivalence of the representations (3.1) and (3.2) follows from Propositions 2.4 and 2.5. The matrix \widehat{C} may be viewed as a constituency matrix for the fluid model (3.2).

If $n \ll \ell_r$, then the behavior of this system may be entirely unnatural since so many constraints have been removed. We show in section 4 that this error can be bounded when considering optimal-control solutions for the fluid model. Related results are obtained for the stochastic model in section 4.2. Such solidarity requires that $n \geq 1$ be chosen sufficiently large, but in many examples this is significantly smaller than ℓ_r .

Our goal remains the same: We wish to minimize, over all feasible state trajectories, the infinite-horizon cost

$$(3.4) \quad \widehat{J}(x) = \int_0^\infty c(\widehat{q}(t; x)) dt, \quad x \in \mathbf{X}.$$

Procedures for *translation* of an optimal allocation \widehat{z}^* to both the original fluid model and to the stochastic model (2.1) are treated in sections 4.1 and 4.2, respectively.

In analogy with (2.17), the workload process for this model is given by

$$\widehat{w}(t; x) = \widehat{\Xi} \widehat{q}(t; x), \quad t \geq 0.$$

For all $1 \leq i \leq n$ we retain the simple constraint

$$(3.5) \quad \frac{d}{dt} \widehat{w}_i(t; x) \geq -(1 - \rho_i), \quad t \geq 0.$$

These constraints are *decoupled* under our assumption that the workload vectors are linearly independent. However, the workload process is also constrained to the set

$$(3.6) \quad \widehat{W} := \{\widehat{\Xi}x : x \in \mathbf{X}\}.$$

The set $\widehat{W} \subseteq \mathbb{R}^n$ is a convex cone since $\mathbf{X} = \mathbb{R}_+^\ell$. In general, $\widehat{W} \not\subseteq \mathbb{R}_+^n$ since elements of a workload vector $w \in \widehat{W}$ may have negative entries.

Two states $x, y \in \mathbf{X}$ are called *exchangeable* if $\widehat{\Xi}(x - y) = \mathbf{0}$. Letting $\widehat{T}^*(x, y)$ denote the optimal time to travel from x to y ,

$$\widehat{T}^*(x, y) = \left(\max_{1 \leq i \leq n} \frac{\langle \xi^i, x - y \rangle}{1 - \rho_i} \right)^+,$$

we see that $\widehat{T}^*(x, y) = \widehat{T}^*(y, x) = 0$ when x and y are exchangeable.

If one is interested in optimal control, then of course one will never stay in a state x if there exists an exchangeable state y with lower cost. Hence an optimal trajectory \widehat{q}^* can always be chosen so that it takes values in

$$\widehat{X} = \{x \in \mathbf{X} : c(x) \leq c(y) \text{ whenever } \widehat{\Xi}x = \widehat{\Xi}y\}.$$

This is an example of *state space collapse* as described in the introduction.

This reasoning leads to the following definitions:

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

$$(3.7) \quad \begin{array}{ll} \min & \gamma \\ \text{subject to} & \gamma \geq \langle c^i, x \rangle, \quad 1 \leq i \leq \ell_c, \\ & \widehat{\Xi}x = w, \\ & x \in \mathbf{X}, \end{array}$$

where $\{c^i\}$ are the components of the cost function given in (2.8). The effective cost is piecewise linear:

$$(3.8) \quad \bar{c}(w) = \max_i \langle \bar{c}^i, w \rangle, \quad w \in \widehat{W},$$

where $\{\bar{c}^i\} \in \mathbb{R}^n$ are the extreme points obtained in the dual of (3.7).

(ii) For any $w \in \widehat{W}$, the *effective state* $\mathcal{X}^*(w)$ is defined to be the vector $x \in \widehat{X}$ that minimizes the linear program (3.7):

$$(3.9) \quad \mathcal{X}^*(w) = \arg \min_{x \in \mathbf{X}} (c(x) : \widehat{\Xi}x = w).$$

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

$$(3.10) \quad \mathcal{P}^*(x) = \mathcal{X}^*(\widehat{\Xi}x).$$

The function \mathcal{X}^* may not be uniquely defined, but it is chosen to be a continuous map from \widehat{W} to \widehat{X} . This is always possible by restricting to basic feasible solutions in (3.7) to obtain a piecewise linear function of x .

Let $\widehat{W}^+ \subset \mathbb{R}^n$ denote the closed, positive cone

$$(3.11) \quad \widehat{W}^+ = \{w \in \widehat{W} : \bar{c}(w) \leq \bar{c}(w') \text{ whenever } w' \geq w, w' \in \widehat{W}\}.$$

The function $\bar{c}: \widehat{W} \rightarrow \mathbb{R}_+$ is called *monotone* if $\widehat{W}^+ = \widehat{W}$ and $\widehat{W} \subseteq \mathbb{R}_+^n$.

Let $\widehat{q}^*(\cdot; x)$ denote an optimal trajectory for the relaxed control problem with initial condition x , and let $\widehat{w}^*(\cdot; x)$ denote the corresponding workload process. By optimality we have the equivalence

$$c(\widehat{q}^*(t; x)) = \bar{c}(\widehat{w}^*(t; x)), \quad t \geq 0.$$

PROPOSITION 3.1. *Suppose that the n th relaxation is stabilizable. Then, the optimal trajectory \widehat{q}^* can be chosen so that for any initial condition $x \in X$,*

- (i) $c(\widehat{q}^*(t; x))$ is decreasing, convex, and piecewise linear,
- (ii) both \widehat{q}^* and \widehat{w}^* are piecewise linear and continuous on $(0, \infty)$.

Proof. The proof of (i) is identical to the result for the original network model (see [36, Proposition 5]).

To see (ii), consider first the workload process. By convexity, $\bar{c}(\widehat{w}^*(t; x))$ can be discontinuous only at $t = 0$. Moreover, we may assume that \widehat{w}^* is linear on each of the open intervals (T_i, T_{i+1}) , $1 \leq i \leq m - 1$, where $\{T_i\}$ denotes the times at which $\frac{d}{dt}c(\widehat{q}^*(t; x))$ is discontinuous, with $T_0 = 0$, $T_m = \infty$.

We now show that, without any loss of generality, the trajectory \widehat{w}^* can be taken to be continuous on $(0, \infty)$. Consider the second time-interval $[T_1, T_2]$. We consider the linear path on this interval given by

$$\widehat{w}^\circ(t) = \widehat{w}^*(T_1-; x) + \frac{t - T_1}{T_2 - T_1} \left[\widehat{w}^*(T_2-; x) - \widehat{w}^*(T_1-; x) \right], \quad T_1 < t < T_2.$$

The identity $\bar{c}(\widehat{w}^\circ(t)) = \bar{c}(\widehat{w}^*(t; x))$ holds on this interval since $\bar{c}(\widehat{w}^*(T_1-; x)) = \bar{c}(\widehat{w}^*(T_1+; x))$.

The trajectory \widehat{w}° is feasible, and we can thus redefine \widehat{w}^* on $(0, T_2)$ so that it is continuous. This procedure can be continued on each interval to form an optimal solution that is continuous on $(0, \infty)$.

To show that \widehat{q}^* can also be taken as continuous, choose $\widehat{q}^*(t; x) = \mathcal{X}^*(\widehat{w}^*(t; x))$, $t > 0$. \square

3.2. One-dimensional workload. The workload process for the relaxed control problem frequently admits an identifiable optimal solution, and in many instances this solution is pointwise optimal.

In the one-dimensional case the matrix $\widehat{\Xi}$ is a row vector, $\widehat{\Xi} = \xi^{1\tau}$. Provided $\rho = \rho_1 < 1$, the minimal draining time is given by

$$\widehat{T}^*(x) = \frac{\max(0, \langle \xi^1, x \rangle)}{1 - \rho_1}, \quad x \in X.$$

The following results follow from linearity of \widehat{T}^* and radial homogeneity of the cost function.

PROPOSITION 3.2. *The following hold for the one-dimensional relaxation for any piecewise linear cost function:*

(i) *The velocity set \widehat{V} is the half space*

$$\widehat{V} = \{v : \langle \xi^1, v \rangle \geq -(1 - \rho)\}.$$

(ii) *The effective cost \bar{c} and the lifting map \mathcal{X}^* are linear functions of w , for $w \geq 0$. Hence, letting $x^* = \mathcal{X}^*(1)$, the following hold for any $w \geq 0$ and any $x \in \mathbf{X}$ satisfying $\langle \xi^1, x \rangle \geq 0$:*

$$\bar{c}(w) = wc(x^*), \quad \mathcal{X}^*(w) = wx^*, \quad \mathcal{P}^*(x) = \langle \xi^1, x \rangle x^*.$$

(iii) *For any $x \in \mathbf{X}$ satisfying $\langle \xi^1, x \rangle \geq 0$, an optimal state trajectory is given by*

$$\widehat{q}^*(t; x) = \mathcal{P}^*(x) \left(\frac{\widehat{T}^*(x) - t}{\widehat{T}^*(x)} \right), \quad 0 < t \leq \widehat{T}^*(x).$$

(iv) *If the initial condition $x \in \mathbf{X}$ satisfies $\langle \xi^1, x \rangle \leq 0$, then an optimal solution is given by $\widehat{q}^*(t; x) = \boldsymbol{\theta}$ for $t > 0$. \square*

Proposition 3.3 shows that the solution in (iii) is pointwise optimal.

PROPOSITION 3.3. *Consider the relaxed control problem with $n = 1$. For any monotone, convex cost function $c: \mathbf{X} \rightarrow \mathbb{R}_+$ and any initial condition, there exists a pointwise optimal allocation.*

Proof. Let $x \in \mathbf{X}$ be given. If $\langle \xi^1, x \rangle \leq 0$, then $\widehat{q}^*(t; x) = \boldsymbol{\theta}$ for all $t > 0$. This is a pointwise optimal solution.

The proof is by comparison when $\langle \xi^1, x \rangle > 0$. Let $x^*(t)$ be the solution to the nonlinear program

$$\begin{aligned} \min \quad & c(y) \\ \text{subject to} \quad & y = x + \widehat{v}t, \\ & \langle \xi^1, \widehat{v} \rangle \geq -(1 - \rho), \\ & y \geq \boldsymbol{\theta}. \end{aligned}$$

Its value, $\widehat{c}^* = c(x^*(t))$, is a lower bound on $c(\widehat{q}(t; x))$ for any feasible state trajectory \widehat{q} since we are optimizing over all states attainable at time t . Moreover, because \widehat{V} is a half-space, the state trajectory $\widehat{q}^*(t; x) = x^*(t)$, $t > 0$, is feasible for the relaxed fluid model. \square

When c is linear, the effective cost has the following specific form:

$$\bar{c}(w) = \left(\frac{c_i^*}{\xi_i^1} \right) w = \left(\min \frac{c_i}{\xi_i^1} : \xi_i^1 > 0 \right) w, \quad w \geq 0,$$

and $x^* = (\xi_{i^*}^1)^{-1} e^{i^*}$. In this case, Proposition 3.2(ii) may be viewed as a generalization of the $c\mu$ -rule [6, 30].

The routing model shown in Figure 2 was used in [29] to illustrate a form of state space collapse for a stochastic model. We assume that the router with service rate μ_3 is fast, so that, in particular, $\mu_3 > \mu_1 + \mu_2$.

The fluid model is given by

$$(3.12) \quad B = \begin{bmatrix} -\mu_1 & 0 & \mu_3 & 0 \\ 0 & -\mu_2 & 0 & \mu_3 \\ 0 & 0 & -\mu_3 & -\mu_3 \end{bmatrix}, \quad \alpha = \begin{bmatrix} 0 \\ 0 \\ \alpha_3 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}.$$




FIG. 2. On the left is shown a simple routing model. At right is the velocity set \mathbf{V} , and its one-dimensional relaxation, projected onto $\{v \in \mathbb{R}^3 : v_3 = 0\}$.

We have four workload vectors,

$$\begin{aligned} \xi^1 &= (\mu_1 + \mu_2)^{-1}(1, 1, 1)^T, & \eta^1 &= (\mu_1 + \mu_2)^{-1}(\mu_1, \mu_2, 0)^T, \\ \xi^2 &= (m_1, 0, 0)^T, & \eta^2 &= e^1, \\ \xi^3 &= (0, m_2, 0)^T, & \eta^3 &= e^2, \\ \xi^4 &= (0, 0, m_3)^T, & \eta^4 &= e^3, \end{aligned}$$

where $m_i = 1/\mu_i$. The vector ξ^1 defines the workload at the two downstream stations, pooled together to form a single resource.

The respective loads are given by $\rho_1 = \alpha_3/(\mu_1 + \mu_2)$, $\rho_2 = \rho_3 = 0$, $\rho_4 = \alpha_3/\mu_3$. The system load is $\rho = \max(\rho_1, \rho_4) = \rho_1$ since we have assumed that $\mu_3 > \mu_1 + \mu_2$. Using the formula given in Theorem 2.2 we can compute the minimum emptying time from an initial condition $x \in \mathbf{X} = \mathbb{R}_+^3$:

$$T^*(x) = \max\left(\frac{1}{1 - \rho_1} \frac{x_1 + x_2 + x_3}{\mu_1 + \mu_2}, \frac{x_1}{\mu_1}, \frac{x_2}{\mu_2}, \frac{1}{1 - \rho_4} \frac{x_3}{\mu_3}\right).$$

Given the expression for ξ^1 we find that the velocity set for the first workload-relaxation is given by

$$\hat{\mathbf{V}} = \{v : v_1 + v_2 + v_3 \geq -(\mu_1 + \mu_2 - \alpha_3)\}.$$

This set is compared to the entire velocity set \mathbf{V} in Figure 2. Although both are defined to be a subset of \mathbb{R}^3 , we can set $q_3 = v_3 \equiv 0$ to obtain the two-dimensional projection shown. We have $\hat{\mathbf{W}} = \mathbb{R}_+$ in the first workload-relaxation, and if the cost is linear, $c(x) = c^T x$, $x \in \mathbf{X}$, then the effective cost is given by

$$c(w) = c_{i_*}(\mu_1 + \mu_2)w, \quad w \in \hat{\mathbf{W}},$$

where $c_{i_*} = \min(c_1, c_2, c_3)$.

3.3. Dimension two. Under certain conditions on the cost we can again be assured of a pointwise optimal solution even when $\hat{\mathbf{V}}$ is not a half-space. We illustrate this in the two-dimensional case where

$$\begin{aligned} \hat{\Xi} &= [\xi^1 \mid \xi^2]^T, \\ \hat{\mathbf{V}} &= \{v : \langle \xi^i, v \rangle \geq -(1 - \rho_i), \quad i = 1, 2\}. \end{aligned}$$

The following result holds for any piecewise linear cost function. Recall the definition of the monotone set given in (3.11).

THEOREM 3.4. *Suppose that $\widehat{W} \subseteq \mathbb{R}_+^2$.*

(i) *When the initial condition satisfies $\widehat{w}(0) \in \widehat{W}^+$, then there exists a pointwise optimal solution.*

(ii) *If the vector $(1 - \rho_1, 1 - \rho_2)^T$ lies in \widehat{W}^+ , then there is a pointwise optimal solution from any initial condition.*

Proof. The proof follows from the rectangular geometry of the set of all states reachable from $\widehat{w}(0) = w$: If w^1 can be reached from w at time t using some allocation, then any $w^2 \in \widehat{W}$ can also be reached, provided $w_i^2 \geq w_i^1$ for each i . Under the conditions imposed in (i), using the greedy policy we have $\widehat{w}^*(t; x) \in \widehat{W}^+$ for all $t > 0$, and $\widehat{w}^*(t; x)$ is pointwise minimal within \widehat{W}^+ in the sense that

$$\widehat{w}_i^*(t; x) \leq \widehat{w}_i(t; x), \quad t \geq 0, \quad i = 1, 2,$$

for any other feasible trajectory \widehat{w} evolving in \widehat{W}^+ , starting at $w = \widehat{\Xi}x$. The result (ii) then follows from (i) since $\widehat{w}_i^*(t; x) \in \widehat{W}^+$ for all $t > 0$ under the assumptions of (ii). \square

Figure 3 shows the structure of the cost function, the set \widehat{W}^+ , and optimal state trajectories for a model that satisfies the assumptions of Theorem 3.4(ii).

Pathwise optimality cannot be expected in general. If the workload dimension is greater than one, and if the cost function c favors starvation of some dynamic bottleneck from some initial condition, then the greedy policy is not time-optimal and hence it cannot be pointwise optimal. Figure 4 illustrates one example with $(1 - \rho_1, 1 - \rho_2) \notin \widehat{W}^+$ and $\widehat{w}(0) \notin \widehat{W}^+$. The initial condition satisfies

$$\frac{d}{dw_2} c(\widehat{w}(0)) < 0.$$

From this initial condition it is advantageous in the short term to allow $\widehat{w}(0+) \in \partial\widehat{W}^+$ since \bar{c} is not monotone. This is the greedy, or myopic, policy, which is not time-optimal in this example. The paths shown minimize the infinite-horizon cost given in




FIG. 3. The figure at left shows a level set of the cost function \bar{c} and the positive cone \widehat{W}^+ on which the cost is monotone. The plot at right shows three optimal state trajectories from varying initial conditions. The darkest region in each figure shows workload vectors w that are not feasible.




FIG. 4. Shown are two optimal state trajectories from two different initial conditions. This example is exactly as before, except that ρ_2 is somewhat larger. A trade-off must be made in this case: An overly greedy decision at time $0+$ will significantly extend the time to empty the network.

(3.4), or equivalently $\hat{J} = \int_0^\infty \bar{c}(\hat{w}(t; x)) dt$. An optimal allocation makes a trade-off between reducing the cost at time $0+$ and preserving a fast draining time for the model, whenever $\hat{w}(0) \notin \hat{W}^+$.

The three-buffer model shown in Figure 5 is described by the linear fluid model with parameters

$$(3.13) \quad B = \begin{bmatrix} -\mu_1 & 0 & 0 \\ \mu_1 & -\mu_2 & 0 \\ 0 & \mu_2 & -\mu_3 \end{bmatrix}, \quad \alpha = \begin{bmatrix} \alpha_1 \\ 0 \\ 0 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}.$$

The load parameters and workload vectors are given by

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

where we have used $\rho = \hat{\Xi}\alpha$, with $\hat{\Xi}$ given in (3.3) with $n = 2$, and $m_i = \mu_i^{-1}$.

Figure 6 shows the optimal solutions for the first and second workload-relaxations. In this numerical example we have taken $\rho_1 = \rho_2 = 9/10$ and $c = (1, 1, 1)^T$. The two plots are very different since the loads at stations one and two are equal.

In Figure 7 the optimal trajectory minimizing (2.6) is compared to the pointwise optimal solution for the two-dimensional relaxation. The triangular region shows the (apparent) error introduced by relaxing the original network optimization problem. We introduce a procedure in Theorem 4.1 below to translate the solution of the relaxed




FIG. 5. A two station scheduling problem




FIG. 6. Optimal cost-curves for the first and second workload-relaxations.

problem to the original network model. This yields precisely the optimal policy in this example.

Figure 1 shows a *pull model* in which four of the buffers are virtual. This example is analyzed in [14] under the assumption that the arrivals are *controlled*. An optimal policy will simultaneously determine sequencing and routing rules at each station and release rules for material to the network. Specific service rate values may be found in Chapter 3 of [14]. The cost c is linear, with a weighting of 10 for deficit and unity weighting at the two other virtual buffers and all real buffers.

Although the model is complex, the effective cost for the second workload-relaxation is very simple: as shown in Figure 8, it is defined by five linear functions $\{\bar{c}^i, i = 1, \dots, 5\}$. Figure 8 shows that the set \hat{W}^+ contains the ray $\{w \in \mathbb{R}_+^2 : \bar{c}^1 w = \bar{c}^2 w\}$ but not much else, since both \bar{c}^1 and \bar{c}^2 have negative components. It follows that pointwise optimal trajectories exist for each initial condition only for a very small set of arrival-rate vectors α . (In this example, arrivals are interpreted as *demand* of material from the network.)

3.4. Higher workload dimension. The two-dimensional case is special because one can always find, for each initial w and each time t , a workload vector $\hat{w}^*(t) \in \hat{W}^+$ that is pointwise minimal and reachable from w at time t . This geometry breaks down in three or more dimensions.

Consider first some positive results.

THEOREM 3.5. *Suppose that $\hat{W} \subseteq \mathbb{R}_+^n$ for the n th workload-relaxation. The following are then equivalent:*

- (i) *A pointwise minimal solution \hat{w}^* exists for any initial condition $x \in X$ and any arrival-rate vector α .*




FIG. 7. The dashed line shows the cost $c(\hat{q}^*(t; x))$ for the optimized two-dimensional workload-relaxation. The actual optimal policy incurs a higher cost, but this error is bounded in ρ .




FIG. 8. On the left is shown a sublevel set of the effective cost for one set of parameters in a two-dimensional relaxation of the network shown in Figure 1. The figure at right shows the set \widehat{W}^+ together with a close-up of the sublevel set of \bar{c} , restricted to \mathbb{R}_+^2 . The workload space \widehat{W} is equal to all of \mathbb{R}^2 in this example.

(ii) For each $w \in \mathbb{R}^n$, the set

$$(3.14) \quad \widehat{W}_w = \{\widehat{w} : \widehat{w} \in \widehat{W}, \widehat{w}_i \geq w_i, \quad i = 1, \dots, n\},$$

contains a pointwise minimal element.

If either of these equivalent conditions hold, then a pointwise minimal trajectory may be expressed,

$$(3.15) \quad \widehat{w}^*(t; x) = [\widehat{\Xi}x - (\mathbf{1} - \rho)t]_+,$$

where $[w]_+$, $w \in \mathbb{R}^n$, is the projection of w onto the set \widehat{W}_w in the standard ℓ_2 norm.

Proof. We first show that the pointwise minimal trajectory is given by (3.15) if (i) holds. Observe that for any t, x the inequality $\widehat{w}^*(t; x) \geq \widehat{\Xi}x - (\mathbf{1} - \rho)t$ holds, and $\widehat{w}^*(t; x) \in \widehat{W}$. Since \widehat{w}^* is minimal, it serves as the projection as claimed.

This implication also shows that (i) \Rightarrow (ii).

Conversely, if (ii) holds, then the trajectory given by $\widehat{w}^\circ(t; x) = [\widehat{\Xi}x - (\mathbf{1} - \rho)t]_+$, $t \geq 0$, is obviously pointwise minimal, and it is a piecewise linear function of t for each initial condition x . We show below that the semigroup property holds:

$$(3.16) \quad \widehat{w}^\circ(t + s; x) = [\widehat{w}^\circ(t; x) - (\mathbf{1} - \rho)s]_+, \quad t, s \geq 0, \quad x \in X.$$

This implies that $\frac{d}{dt}\widehat{w}^\circ(t; x) \geq -(\mathbf{1} - \rho)$ for all t , so that this trajectory is feasible for the relaxed fluid model, and hence (i) holds with $\widehat{w}^* = \widehat{w}^\circ$.

To establish (3.16), fix $s, t > 0$, let $T_1 = s + t$, and consider for comparison the trajectory

$$\widehat{w}(T; x) = \widehat{\Xi}x + \frac{T}{T_1}[\widehat{w}^\circ(t + s; x) - \widehat{\Xi}x], \quad 0 \leq T \leq T_1.$$

This is feasible, and by minimality of $\widehat{w}^\circ(t; x)$ we have $\widehat{w}(t; x) \geq \widehat{w}^\circ(t; x)$. The following bounds then follow:

$$\begin{aligned} \widehat{w}^\circ(t + s; x) = \widehat{w}(T_1; x) &= [\widehat{w}(t; x) + (\widehat{w}(t + s; x) - \widehat{w}(t; x))]_+ \\ &\geq [\widehat{w}(t; x) - (\mathbf{1} - \rho)s]_+ \\ &\geq [\widehat{w}^\circ(t; x) - (\mathbf{1} - \rho)s]_+. \end{aligned}$$

To obtain an inequality in the reverse direction, note that $\widehat{w}^\circ(t; x) \geq \widehat{\Xi}x - (\mathbf{1} - \rho)t$, which implies that

$$\widehat{w}^\circ(t + s; x) := [\widehat{\Xi}x - (\mathbf{1} - \rho)(t + s)]_+ \leq [\widehat{w}^\circ(t; x) - (\mathbf{1} - \rho)s]_+.$$

We therefore obtain (3.16). \square

Under these conditions there is some hope in finding a pointwise optimal solution to the relaxed optimal control problem.

COROLLARY 3.6. *Suppose that*

- (i) *the effective cost \bar{c} is monotone and*
- (ii) *a pointwise minimal solution \widehat{w}^* exists for the n th workload-relaxation, for any initial condition $x \in \mathsf{X}$.*

Then for any $x \in \mathsf{X}$ there is a pointwise optimal solution for the n th workload-relaxation. \square

Assumption (ii) fails in general. Consider the three-station network shown in Figure 9 (see [29, sections 6 and 7] for related examples of RBM networks). The arrival rates α_1, α_6 are equal, and all service rates are equal to unity. For any x , the vector $\widehat{w} = \widehat{\Xi}x \in \mathbb{R}^3$ can be written

$$\widehat{w}_1 = x_1 + x_2 + x_4 + x_6, \quad \widehat{w}_2 = x_1 + x_3 + x_4 + x_6, \quad \widehat{w}_3 = x_6 + x_1 + x_3 + x_5.$$

For example, $\widehat{w}^3 := \widehat{\Xi}e^3 = [0, 1, 1]^T$, and $\widehat{w}^4 := \widehat{\Xi}e^4 = [1, 1, 0]^T$. The two states $\{e^3, e^4\}$ are not exchangeable for a three-dimensional relaxation since the workload vectors $\widehat{w}^3, \widehat{w}^4$ are different.

For simplicity consider the arrival-free model where $\alpha_1 = \alpha_6 = 0$ so that $\rho = 0$. The initial condition $x = e^3 + e^4$ has corresponding workload $\widehat{w}(0) = \widehat{\Xi}x = (1, 2, 1)^T$. From this initial condition it is possible to reach either e^3 or e^4 in exactly one second. Any minimal workload vector \widehat{w}^* must then satisfy $\widehat{w}^*(t; x) \leq \widehat{w}^3$ and $\widehat{w}^*(t; x) \leq \widehat{w}^4$ at $t = 1$, which implies that $\widehat{w}^*(1; x) \leq (0, 1, 0)^T$.

The only vector in \widehat{W} satisfying this inequality is $w = (0, 0, 0)^T$. However, this state is not reachable in one second since the minimal draining time is $W^*(x) = T^*(x) = \max(\widehat{w}_1(0), \widehat{w}_2(0), \widehat{w}_3(0)) = 2$.

We now investigate the structure of pointwise optimal solutions under the conditions of Corollary 3.6.

The i th pooled-resource is said to be *satiated* at state x provided there exists $v \in \widehat{V}$ satisfying $\langle \xi^i, v \rangle = -(1 - \rho_i)$, and $v_i \geq 0$ whenever $x_i = 0$. A resource is said to be satiated if it is a component of a satiated pooled-resource.

Consider any $x \in \mathsf{X}$, and suppose $y \in \mathsf{X}$ with $\langle \xi^k, x \rangle > \langle \xi^k, y \rangle$ for some $1 \leq k \leq n$. Then the optimal time to travel from x to y is nonzero:

$$\widehat{T}^*(x, y) = \max_{1 \leq j \leq n} \frac{\langle \xi^j, x - y \rangle}{1 - \rho_j} > 0.$$




FIG. 9. A three-station network.

With $v = (y-x)/\widehat{T}^*(x, y) \in \widehat{V}$, the trajectory below is both feasible and time-optimal:

$$\widehat{q}(t; x) = x + vt, \quad 0 \leq t \leq \widehat{T}^*(x, y).$$

Moreover, simple dynamic programming arguments ensure that

$$\frac{d}{dt} \widehat{T}^*(\widehat{q}(t; x), y) = -1, \quad 0 < t < \widehat{T}^*(x, y).$$

Hence, whenever i is a maximizer, so that

$$\widehat{T}^*(x, y) = \frac{\langle \xi^i, x - y \rangle}{1 - \rho_i},$$

we must have $\langle \xi^i, v \rangle = -(1 - \rho_i)$. This implies that pooled-resource i is satiated by x and proves the following.

PROPOSITION 3.7. *Suppose that $\rho < 1$. Then $\widehat{T}^*(x, y) < \infty$, $x, y \in X$, and if $\widehat{T}^*(x, y) > 0$, then*

$$\begin{aligned} \widehat{T}^*(x, y) &= \max \left\{ \frac{\langle \xi^j, x - y \rangle}{1 - \rho_j} : 1 \leq j \leq n \right\} \\ &= \max \left\{ \frac{\langle \xi^j, x - y \rangle}{1 - \rho_j} : j \text{ is satiated by } x \right\}. \quad \square \end{aligned}$$

Satiated resources play a role analogous to dynamic bottlenecks in the construction of a time-optimal trajectory. The following result is the analogue of Theorem 2.6. It is an easy corollary to Proposition 3.7.

THEOREM 3.8. *Suppose that $\rho < 1$, and that the n th relaxation satisfies $\widehat{W} \subseteq \mathbb{R}_+^n$. Let \widehat{q} be any solution to the n th workload-relaxation, starting at $x \in X$, and let $\widehat{w}(t; x) = \widehat{\Xi} \widehat{q}(t; x)$, $t \geq 0$. We then have the following:*

(i) *If \widehat{w} is pointwise minimal, then each satiated pooled-resource is working at capacity for each $0 < t < \infty$. That is, if pooled-resource i is satiated at time t , then*

$$\frac{d}{dt} \widehat{w}_i(t) = -(1 - \rho_i).$$

(ii) *If each satiated pooled-resource works at capacity for all t , then the resulting workload trajectory \widehat{w} is pointwise minimal. \square*

3.5. Variability and continuity. We close this section with some continuity properties for pointwise minimal solutions. Our interest lies in the fluid model with *exogenous disturbance*, defined by

$$(3.17) \quad \widehat{q}(t; x) = B\widehat{z}(t; x) + \alpha t + d_0(t), \quad t \geq 0, \quad \widehat{q}(0; x) = x.$$

We assume as in (3.2) that the allocation is subject to the linear constraints

$$(3.18) \quad \widehat{C}[\widehat{z}(t; x) - \widehat{z}(s; x)] \leq [t - s]\mathbf{1}, \quad 0 \leq s \leq t,$$

where $\widehat{C} = -\widehat{\Xi}B$ is defined below (3.2), and we assume throughout that the disturbance d_0 is of bounded variation.

Letting $\widehat{w}(t; x) := \widehat{\Xi}\widehat{q}(t; x)$, $d(t) = \widehat{\Xi}d_0(t)$, we obtain the corresponding workload model

$$(3.19) \quad \widehat{w}(t; x) = \widehat{\Xi}x - (\mathbf{1} - \boldsymbol{\rho})t + \boldsymbol{\iota}(t) + d(t), \quad t \geq 0,$$

where $\boldsymbol{\iota}(t) := t\mathbf{1} - \widehat{C}\widehat{z}(t; x)$, $t \geq 0$. The idleness process $\boldsymbol{\iota}$ is nonnegative with nondecreasing components, and \widehat{w} evolves in \widehat{W} .

Rather than define \widehat{w} through (3.17), for the purposes of optimization we may restrict attention to the simpler model (3.19). Given the current workload-value $\widehat{w} = \widehat{w}^*(t; x)$ we take $\widehat{z}^*(t; x)$ to be any optimizer of the linear program

$$(3.20) \quad \begin{aligned} & \min \quad \gamma \\ & \text{subject to} \quad \gamma \geq \langle \mathbf{c}^i, \mathbf{y} \rangle, \quad 1 \leq i \leq \ell_c, \\ & \quad \quad \quad \mathbf{y} = \mathbf{x} + B\mathbf{z} + \alpha t + d_0(t), \\ & \quad \quad \quad \widehat{\Xi}\mathbf{y} = \widehat{w}, \\ & \quad \quad \quad \mathbf{y} \in \mathbf{X}. \end{aligned}$$

It follows from the definitions that the optimizer \widehat{z}^* satisfies the constraints given in (3.18).

If $\mathbf{d} \equiv \boldsymbol{\theta}$, then (3.19) is the linear workload model considered earlier. In this case, it follows from Theorem 3.5 and Assumption M below that (3.19) admits a pointwise minimal solution for any value of $\boldsymbol{\rho}$. This is generalized to arbitrary disturbances in Theorem 3.10.

Assumption M.

- (i) The workload vectors for the n th relaxation are linearly independent and satisfy

$$\boldsymbol{\xi}^i \in \mathbb{R}_+^\ell \text{ for } 1 \leq i \leq n.$$

- (ii) For each $w \in \mathbb{R}^n$ the set \widehat{W}_w defined in (3.14) contains a pointwise minimal element denoted $[w]_+$.

Although the semigroup property (3.16) does not hold in general for a model with disturbances, we always have the lower bound.

LEMMA 3.9. *Under Assumption M, if \widehat{w} is a feasible state trajectory for the n th relaxation, then*

$$\widehat{w}(t; x) \geq [\widehat{w}(s; x) - (\mathbf{1} - \boldsymbol{\rho})(t - s) + d(t) - d(s)]_+, \quad t \geq s \geq 0, \quad x \in \mathbf{X}.$$

Proof. The lower bound $\widehat{w}(t; x) \geq (\widehat{w}(s; x) - (\mathbf{1} - \boldsymbol{\rho})(t - s) + d(t) - d(s))$ holds since the idleness process is nondecreasing. Hence the result follows from the definition of the projection, combined with Assumption M, which asserts that the projection can be taken to be pointwise minimal. \square

Theorem 3.10 establishes existence of minimal solutions and some strong robustness properties. This existence question is closely related to the *generalized Skorokhod problem* [21, 26, 2, 15, 16, 18]. These results will facilitate the treatment of stochastic models in section 4.

THEOREM 3.10. *Under Assumption M, for any given disturbance \mathbf{d} of bounded variation, the model (3.19) admits a solution \widehat{w}^* that is pointwise minimal. For two disturbances $(\mathbf{d}^1, \mathbf{d}^2)$ the respective minimal solutions $(\widehat{w}^{*1}, \widehat{w}^{*2})$ satisfy the following:*

- (i) *Provided $d_0^1(t) \leq d_0^2(t)$, $t \geq 0$,*

$$\widehat{w}^{*2}(t; x) \leq \widehat{w}^{*1}(t; x) + d^2(t) - d^1(t), \quad t \geq 0, \quad x \in \mathbf{X}.$$

(ii) Suppose that $d_0^1(t) = d_0^2(t) - \varepsilon_0(t)$, $t \geq 0$, with $\varepsilon_0(\cdot)$ a nonnegative and nondecreasing function from \mathbb{R}_+ to \mathbb{R}_+^ℓ . Then,

$$\widehat{w}^{*1}(t; x) \leq \widehat{w}^{*2}(t; x), \quad t \geq 0, x \in \mathbf{X}.$$

(iii) For arbitrary disturbances $\mathbf{d}_0^1, \mathbf{d}_0^2$,

$$\widehat{w}^{*1}(t; x) \leq \widehat{w}^{*2}(t; x) + |\mathbf{d}^2 - \mathbf{d}^1|_\infty^t - [d^2(t) - d^1(t)], \quad t \geq 0, x \in \mathbf{X},$$

where $(|\mathbf{f}|_\infty^t)_i := \sup_{0 \leq s \leq t} |f_i(s)|$ for any function $\mathbf{f}: \mathbb{R}_+ \rightarrow \mathbb{R}^n$.

Proof. We first establish the three properties, given that minimal solutions exist.

To prove (i), observe that if the optimal allocation $\widehat{\mathbf{z}}^{*1}$ for the first system is applied to the second, then we have for all $t \geq 0$

$$(3.21) \quad \begin{aligned} \widehat{q}^{*1}(t; x) &= x + B\widehat{\mathbf{z}}^{*1}(t) + \alpha t + d_0^1(t), \\ \text{and } \widehat{q}^2(t; x) &= x + B\widehat{\mathbf{z}}^{*1}(t) + \alpha t + d_0^2(t) \geq \widehat{q}^{*1}(t; x) \geq \boldsymbol{\theta}. \end{aligned}$$

Hence $\widehat{\mathbf{z}}^{*1}$ is feasible for the second disturbance, and consequently $\widehat{w}^{*2}(t; x) \leq \widehat{w}^2(t; x)$, with $\widehat{w}^2(t; x) := \widehat{\Xi}\widehat{q}^2(t; x)$, by the assumed existence of a minimal process $\widehat{\mathbf{w}}^{*2}$. Moreover, (3.21) implies that $\widehat{w}^2(t; x) = \widehat{w}^{*1}(t; x) + d^2(t) - d^1(t)$, $t \geq 0$, which gives (i).

The proof of (ii) is similar: Define $\varepsilon(t) := \widehat{\Xi}\varepsilon_0(t)$, $t \geq 0$. Under Assumption M and the conditions imposed in (ii), this is nonnegative and nondecreasing. Let $\widehat{i}^{*i}(t) = \mathbf{1}t - \widehat{C}\widehat{\mathbf{z}}^{*i}(t)$, $i = 1, 2$, denote the optimal idleness, and set $\widehat{i}^1(t) = \widehat{i}^{*2} + \varepsilon(t)$, $t \geq 0$. We have $\frac{d}{dt}\widehat{i}^1(t) \geq \boldsymbol{\theta}$, and we also have under this policy, applied to the first model, $\widehat{w}^1(t; x) = \widehat{w}^{*2}(t; x)$. This combined with minimality of $\widehat{\mathbf{w}}^{*1}$ proves (ii).

To prove (iii) let \mathbf{d}_0^3 denote the disturbance $d_0^3(t) = d_0^1(t) + |\mathbf{d}_0^2 - \mathbf{d}_0^1|_\infty^t$, and let $\widehat{\mathbf{w}}^{*3}$ denote the associated minimal solution. We have $d_0^3(t) \geq d_0^2(t)$, and $\varepsilon_0(t) := d_0^3(t) - d_0^1(t)$ is nonnegative and nondecreasing. Consequently, for any $t \geq 0$, $x \in \mathbf{X}$, we have

$$\begin{aligned} \widehat{w}^{*3}(t; x) &\leq \widehat{w}^{*2}(t; x) + |\mathbf{d}^2 - \mathbf{d}^1|_\infty^t + [d^1(t) - d^2(t)] && \text{from (i),} \\ \widehat{w}^{*1}(t; x) &\leq \widehat{w}^{*3}(t; x) && \text{from (ii).} \end{aligned}$$

Combining these bounds gives (iii).

We now establish existence. Consider first the special case in which all of the disturbances are continuous and piecewise linear. In this case we may construct a pointwise minimal trajectory $\widehat{\mathbf{w}}^*$ inductively by adapting the construction used in Theorem 3.5. Set $\widehat{w}^*(0; x) = \widehat{\Xi}x$, and

$$\widehat{w}^*(T_k + t; x) = [\widehat{w}^*(T_k; x) - (\mathbf{1} - \boldsymbol{\rho})t + m_k t]_+, \quad 0 < t < T_{k+1} - T_k, \quad k \geq 0,$$

where $\{T_i\}$ are the times at which the slope of \mathbf{d} changes, and m_k denotes the slope of \mathbf{d} on the interval $[T_k, T_{k+1}]$. An application of Theorem 3.5 shows that this is the desired minimal solution on $[T_k, T_{k+1}]$ with initial condition $\widehat{w} = \widehat{w}^*(T_k; x)$, and by induction it follows that $\widehat{\mathbf{w}}^*$ is pointwise minimal.

For an arbitrary disturbance \mathbf{d} of bounded variation we can construct a sequence of piecewise linear functions $\{\mathbf{d}^k\}$ such that $d^k(t) \downarrow d(t)$, $k \rightarrow \infty$. We let $\{\widehat{\mathbf{w}}^{*k}\}$ denote the respective optimal solutions and set $\widehat{w}^*(t; x) = \liminf_{k \rightarrow \infty} \widehat{w}^{*k}(t; x)$ for all t, x . Using property (i) for the $\{\widehat{\mathbf{w}}^{*k}\}$ we deduce that $\widehat{\mathbf{w}}^*$ is the desired pointwise minimal solution. \square

We see that it is frequently possible to compute a pointwise optimal trajectory $\widehat{\mathbf{q}}^*$ for the relaxed control problem, with or without disturbances. What does this

then tell us about the original model of interest? The sharpest results are obtained by examining a model in heavy traffic, with $\rho \sim 1$.

4. Networks in heavy traffic. We consider here a sequence of networks, indexed by an integer $r \geq 1$, for which $\rho^r \uparrow 1$ as $r \rightarrow \infty$. It is in this heavily loaded regime that the time-scale separation developed in the previous section is most evident in the (unrelaxed) network model.

We assume that B and C are independent of r . Two arrival-rate vectors α^1, α^∞ are given, and for arbitrary $r \geq 1$ we set

$$(4.1) \quad \alpha^r := \alpha^\infty - \frac{1}{r}(\alpha^\infty - \alpha^1).$$

We impose the following assumptions throughout this section.

Assumption H.

(i) The model with arrival-rate vector α^1 is stabilizable. In particular,

$$\rho^1 := W^*(\alpha^1) < 1.$$

(ii) The arrival-rate vector α^∞ satisfies $\alpha^1 \leq \alpha^\infty$ and

$$\rho^\infty := W^*(\alpha^\infty) = 1.$$

We let $\mathcal{I}_b = \{i : \langle \xi^i, \alpha^\infty \rangle = 1\}$ denote the index set of bottleneck stations for the model with arrival rate α^∞ . By reordering, we can assume, without loss of generality, that $\mathcal{I}_b = \{1, \dots, \ell_b\}$ for some integer $\ell_b \geq 1$.

The choice of a perturbation in the arrival stream is for the sake of convenience since we can then take a fixed set of workload vectors. If we assume that \mathbb{V}_r is a general, convergent sequence of polyhedra, then the theory below remains essentially unchanged.

Throughout this section we consider the n th workload-relaxation with $n = \ell_b$.

4.1. Fluid models. The r th network is defined on a fluid scale by

$$(4.2) \quad \frac{d}{dt}q(t; x) = B\zeta(t; x) + \alpha^r t, \quad t \geq 0.$$

We let \mathbb{V}_r denote the corresponding velocity space so that $\frac{d}{dt}q(t; x) \in \mathbb{V}_r$ for all t, x, r .

The following bound on ρ^r shows that this model is stabilizable for finite $r \geq 1$. The inequality is obtained using convexity of W^* :

$$(4.3) \quad \begin{aligned} \rho^r = W^*(\alpha^r) &= W^*\left(\left(1 - \frac{1}{r}\right)\alpha^\infty + \frac{1}{r}\alpha^1\right) \\ &\leq \left(1 - \frac{1}{r}\right)\rho^\infty + \frac{1}{r}\rho^1 = 1 - \frac{1}{r}(1 - \rho^1) < 1. \end{aligned}$$

For finite r we have $\rho_i^r = 1 - r^{-1}\langle \xi^i, \alpha^\infty - \alpha^1 \rangle$, $i \in \mathcal{I}_b$.

Theorem 4.1 shows that little is lost when considering the ℓ_b th relaxation. Let J^* , \widehat{J}^* denote the value functions for the infinite-horizon optimal control problems defined in (2.3), (3.4), respectively. We always have

$$\widehat{J}^*(x) \leq J^*(x), \quad x \in \mathcal{X}.$$

We obtain a bound in the reverse direction in this section. The analysis is simplest when optimal trajectories are uniquely defined.

Assumption U.

- (i) The linear program (3.9) that defines the effective state $\mathcal{X}^*(w)$ has a unique solution for each $w \in \widehat{W}$.
- (ii) For all $r \geq 1$ sufficiently large and each $T > 0$, $x \in \mathbf{X}$, the ℓ_b th workload-relaxation admits a solution \widehat{q}^{r*} that minimizes the total cost (2.6), and this solution is *unique*.

Consider for example the one-dimensional relaxation of the simple routing model shown in Figure 2. Assume that the cost is linear, so that $c(x) = \langle c, x \rangle$, with $c \in \mathbb{R}_+^3$. If $c_3 \geq c_2 > c_1$, then the above conditions hold. The greedy priority policy that prefers routing to buffer 1, whenever buffer 2 is nonempty, is the unique (pointwise) optimal solution.

Note that Assumption U(i) implies (ii) under Assumption M since in this case $\widehat{q}^*(t; x) = \mathcal{X}^*(\widehat{w}^*(t; x))$, and the pointwise minimal solution \widehat{w}^* is always uniquely defined when it exists.

Applying (3.5) and the form of the rate vector given in (4.1), we find that the constraints on the workload relaxation may be expressed as

$$\frac{d}{dt} \widehat{w}_i(t; x) \geq -\frac{1}{r} \delta_i, \quad 1 \leq i \leq \ell_b, \quad r \geq 1, \quad t > 0,$$

where $\delta_i = \langle \xi^i, \alpha^\infty - \alpha^1 \rangle$. Letting $\widehat{w}^{1*}, \widehat{J}^{1*}$ denote the optimal trajectory and value function when $r = 1$, it follows that for any $r \geq 1$ the optimal solution is given by

$$(4.4) \quad \begin{aligned} \widehat{w}_i^*(t; x) &= \widehat{w}_i^{1*}(t/r; x), & 1 \leq i \leq \ell_b, \quad t > 0, \\ \widehat{J}^*(x) &= r \widehat{J}^{1*}(x), & x \in \mathbf{X}. \end{aligned}$$

We define a policy for the unrelaxed model as follows. Applying Proposition 3.1 we are assured of the existence of a piecewise linear, optimal solution to the relaxed control problem, which we denote $[\widehat{q}^*(t; x), \widehat{\zeta}^*(t; x)]$. The allocation rate $\zeta(t; x)$ for the unrelaxed model is defined to be a function of $[\widehat{q}^*(t; x), \widehat{\zeta}^*(t; x), q(t; x)]$ for any initial condition x and any $t \geq 0$. Let $\mathcal{I}_c(x) = \{i : c(x) = \langle c^i, x \rangle\}$, and given the current states $y = q(t; x)$, $y^* = \widehat{q}^*(t; x)$, let $\zeta(t; x)$ be the optimizing value of the variable ζ in the linear program

$$(4.5) \quad \begin{aligned} \min \quad & \gamma \\ \text{subject to} \quad & \gamma \geq \langle c^i, B\zeta \rangle, \quad i \in \mathcal{I}_c(y), \\ & C\zeta \leq \mathbf{1}, \\ & \zeta \geq \boldsymbol{\theta}, \\ & (B\zeta + \alpha^r)_i \geq 0 \quad \text{if } y_i = 0, \\ & \langle \xi^i, (B\zeta + \alpha^r) \rangle \leq \langle \xi^i, (B\widehat{\zeta}^* + \alpha^r) \rangle, \quad \text{whenever } i \leq \ell_b, \\ & \quad \text{and } \langle \xi^i, y \rangle = \langle \xi^i, y^* \rangle. \end{aligned}$$

The last constraint ensures that $w_i(t; x) \leq \widehat{w}_i^*(t; x)$ for all $i \leq \ell_b$ and all t .

Assume that $q(t; x)$ is the resulting state trajectory using this policy for all t , and set

$$e^r(t; x) = q(t; x) - \widehat{q}^*(t; x), \quad t > 0, \quad \underline{T}_{r,o}(x) = \min\{t : e^r(t; x) = \boldsymbol{\theta}\}.$$

The following result provides uniform bounds on $\underline{T}_{r,o}$ and shows that this first hitting time is in fact a *coupling time*. It is possible to relax the uniqueness assumption in

Theorem 4.1, but one must redefine $\underline{T}_{r_0}(x)$ as the first hitting time to *some* optimal $\widehat{q}^*(t; x)$. A proof is provided in Appendix A.

THEOREM 4.1. *Under Assumptions H and U, the following hold for the state trajectory \mathbf{q} for any initial condition $x \in \mathbf{X}$:*

- (i) $w(t; x) \leq \widehat{w}^*(t; x)$, $t \geq 0$, where the inequality is interpreted componentwise.
- (ii) The time \underline{T}_{r_0} is uniformly bounded in r : For some $b_0 < \infty$,

$$\underline{T}_{r_0}(x) \leq b_0 \|e^r(0+; x)\|, \quad x \in \mathbf{X}, \quad r \geq 1.$$

- (iii) $q(t; x) = \widehat{q}^*(t; x)$ for all $t \geq \underline{T}_{r_0}(x)$.
- (iv) There is a constant b_1 such that

$$\begin{aligned} J(x) &:= \int_0^\infty c(q(t; x)) \leq (1 + b_1/r) \widehat{J}^*(x) \\ &\leq (1 + b_1/r) J^*(x), \quad x \in \mathbf{X}, \quad r \geq 1. \end{aligned}$$

- (v) Suppose that $\widehat{\mathbf{q}}^*$ is a pointwise optimal solution. Then

$$c(q(t; x)) = \underline{c}^*(t; x), \quad t \geq \underline{T}_{r_0}(x),$$

where \underline{c}^* is given in (2.9). □

4.2. Stochastic models. Although the workload-relaxation is in general a significant distortion of the original model, we have seen in Theorem 4.1 that this is negligible when the model is in heavy traffic. The workload constraints overwhelm all other constraints on the velocity vector field. In this section we establish similar solidarity for the stochastic model.

To obtain any such solidarity we must control modelling error, and we must understand when if ever a user can benefit from statistical information. Consider a $G/G/2$ queue, where the two servers are constrained so that only one can work at any given time-instance. The fluid model is given by the one-dimensional model

$$(4.6) \quad \frac{d}{dt} q(t; x) = -\mu_1 \zeta_1(t; x) - \mu_2 \zeta_2(t; x) + \alpha,$$

with the linear control constraint $\zeta_1 + \zeta_2 \leq 1$. This can be viewed as an idealized two-armed bandit, where α is the rate at which a gambler is paying the casino, and $\mu_i \zeta_i$ is his rate of return on using the i th arm. The casino's reward at time t is a linear function of $q(t)$. If $\mu_1 = \mu_2 > \alpha$, then obviously any nonidling policy is optimal, from the gambler's point of view, for any monotone cost function.

For the stochastic model, however, the particular allocation chosen can have great impact since variability of service rates determines the steady-state queue length. For a priority policy in which server i is used exclusively, we obtain in steady-state an approximation of the form, for $\rho \sim 1$,

$$\mathbb{E}[Q(t)] = \frac{1}{2} \frac{\gamma^2}{1 - \rho} + O(1).$$

The infinite-horizon optimal policy is precisely the priority policy that chooses the server with the smallest variability parameter γ^2 .

This example is special because the optimal fluid policy is *not unique*. Typically, the optimal control problem for the fluid model may be solved uniquely since linear

programs generically have unique solutions. If this is the case, then we have fewer opportunities to successfully gamble.

In Theorem 4.3 we impose uniqueness through Assumption U, and an assumption that B is full-rank with $\ell \geq \ell_u$, so that \mathbf{Z} is essentially determined by \mathbf{Q} . The latter assumption may be relaxed considerably by expanding the state space.

Take, for example, the routing model in which B is the 3×4 matrix given in (3.12). Consider the associated four-dimensional network model \mathbf{Q}^a on $\mathbf{X} := \mathbb{R}_+^4$, in which the fourth component is the total-idleness at buffer 1, given by $Q_4^a(t; x) = t - Z_1(t; x)$, $t \geq 0$. The associated matrix B^a is invertible, as seen by the explicit form

$$B^a = \begin{bmatrix} -\mu_1 & 0 & \mu_3 & 0 \\ 0 & -\mu_2 & 0 & \mu_3 \\ 0 & 0 & -\mu_3 & -\mu_3 \\ -1 & 0 & 0 & 0 \end{bmatrix}, \quad \alpha^a = \begin{bmatrix} 0 \\ 0 \\ \alpha_3 \\ 1 \end{bmatrix}.$$

If the cost function on \mathbf{Q}^a is assumed linear, with $c_1 < c_2 < c_3$ and $c_4 > 0$, then Assumption U holds for the four-dimensional model.

For any network model one may augment the state space to include total-idleness, as well as total-allocation values. The cost may be similarly augmented to reflect the desire to maximize utilization of some resources, while minimizing utilization of others. The augmented model will satisfy assumptions (i)–(iii) of Theorem 4.3 for a very general class of network models and cost criteria.

How do we choose the allocation \mathbf{Z} to maintain solidarity with an ideal fluid solution $[\hat{\mathbf{q}}^*, \hat{\mathbf{z}}^*]$? There are three issues that must be addressed:

(i) Suppose that for a given state x , a state $x^* \in \mathbf{X}$ is chosen as a target, with $\hat{\mathbf{e}}x^* \in \hat{\mathbf{W}}^+$. For the fluid model, even if the buffers are empty, an associated resource may be required to work at full capacity. This is not feasible for the discrete model: if a resource finds no work available, then it cannot work. This may be disastrous if the resource is a dynamic bottleneck since any idle time will rule out time-optimality.

(ii) To ensure feasibility we can impose small safety stocks, a well-motivated and standard technique in policy synthesis for manufacturing models [13, 20, 36]. We must ensure that these safety-stock levels can be maintained through a modification of the fluid-allocation without introducing idleness.

(iii) To ensure success we require bounds on the variability of the stochastic processes $(\mathbf{A}, \mathbf{R}, \mathbf{S})$ used in the stochastic model.

To simplify the statements of our assumptions we henceforth assume that the stochastic model (2.1) has the following specific form: For each $1 \leq i \leq \ell$ and $t \geq 0$,

$$Q_i^r(t; x) = x - \sum_{j=1}^{\ell_u} S_{ij}(Z_j(t)) + \sum_{j=1}^{\ell_u} R_{ij}(Z_j(t)) + A_i(t(1 - r^{-1}\delta_i^\alpha)),$$

where $\delta_i^\alpha := (\alpha_i^\infty - \alpha_i^1)/\alpha_i^\infty$ for each i , and the arrival-rate vectors α^1, α^∞ satisfy Assumption H. We assume that the stochastic model is consistent with the fluid model, in the sense that (2.4) holds with $\alpha = \alpha^\infty$. In particular, if $\alpha_i^\infty = 0$, then the process \mathbf{A}_i is null. Assumption S formalizes our remaining probabilistic assumptions. Under this condition we can devise a policy that tracks any fluid idealization and simultaneously ensures that critical resources do not risk starvation.

Assumption S. For all $1 \leq i \leq \ell$ and $1 \leq k \leq \ell_u$, each of the stochastic processes $\{\mathbf{A}_i, \mathbf{R}_{ik}, \mathbf{S}_{ik}, t \geq 0\}$ is either null or is an undelayed renewal process whose increment process possesses a moment generating function that is bounded in a neighborhood

of the origin. The stochastic processes $\{\mathbf{A}, \mathbf{R}, \mathbf{S}\}$ are adapted to a given filtration $\{\mathcal{H}_t : t \geq 0\}$.

We continue to assume that the allocation process \mathbf{Z} satisfies the constraints (2.2), and we assume that any allocation \mathbf{Z} is progressively measurable in the sense that

$$\sigma(Q^r(s), Z^r(s) : s \leq t) \subset \mathcal{H}_t, \quad t \geq 0.$$

A relaxed model $[\widehat{\mathbf{Q}}, \widehat{\mathbf{Z}}]$ is defined in analogy with (3.2), in which the allocation constraint is relaxed to

$$(4.7) \quad \widehat{C}[\widehat{Z}(t; x) - \widehat{Z}(s; x)] \leq [t - s]\mathbf{1}, \quad \widehat{C} := -\widehat{\Xi}^T B.$$

This is of course subject to the additional constraint that $\widehat{Q}(t; x)$ evolves in $\mathbf{X} := \mathbb{R}_+^\ell$. We assume that $\widehat{\mathbf{Z}}$ is of bounded variation, but unlike \mathbf{Z} , it is not subject to any statistical constraints.

For any feasible pair $[\widehat{\mathbf{Q}}, \widehat{\mathbf{Z}}]$ we define the *pseudodisturbance* \mathbf{d}_0 through the equation

$$(4.8) \quad \widehat{Q}(t; x) = x + B\widehat{Z}(t; x) + \alpha^r t + d_0(t), \quad t \geq 0,$$

and we let $d(t) = \widehat{\Xi}d_0(t)$, $t \geq 0$. The associated workload process may be expressed in terms of \mathbf{d} as follows: first define the idleness process by $\widehat{I}(t; x) := t\mathbf{1} - \widehat{C}\widehat{Z}(t; x)$, $t \geq 0$. This is vector-valued, and (4.7) implies that its components are nonnegative and nondecreasing. We then write

$$\widehat{W}(t; x) := \widehat{\Xi}\widehat{Q}(t; x) = -(\mathbf{1} - \rho^r)t + \widehat{I}(t; x) + d(t).$$

We consider below the optimal solution $[\widehat{\mathbf{q}}^*, \widehat{\mathbf{z}}^*]$ to the ℓ_b th fluid-model relaxation (3.17) with respect to the (random) pseudodisturbance \mathbf{d}_0 . This of course depends upon $\widehat{\mathbf{Z}}$. These processes are used for comparison to obtain performance bounds. For example, under the conditions of Theorem 3.10 we obviously have the absolute lower bound, $\widehat{W}(t; x) \geq \widehat{w}^*(t; x) := \widehat{\Xi}\widehat{q}^*(t; x)$, $t \geq 0$. Perhaps surprisingly, the policies considered below almost achieve this lower bound, uniformly for the time-horizons considered.

4.3. Sensitivity and optimality. In the development that follows we construct a trajectory $[\mathbf{Q}^{r^\circ}, \mathbf{Z}^{r^\circ}]$ by attempting to mimic the flow of the optimal fluid trajectory. We begin with a list of desirable properties that $[\mathbf{Q}^{r^\circ}, \mathbf{Z}^{r^\circ}]$ should satisfy. In Theorem 4.3 we show that these general properties imply a strong form of approximate optimality.

Following this we provide a constructive procedure for policy synthesis to attain these properties. This requires some assumptions on the model that we illustrate first in one dimension in section 4.4 and then for general models in section 4.5.

The following result is central to all of the remaining analysis in this section and is essentially our only motivation for Assumption S. A proof may be found in Appendix B.

If $(\mathbf{X}, \mathbf{Y}) = \{(X_r, Y_r) : r \geq 1\}$ is a sequence of random variables, we write $\mathbf{X} \leq O(\mathbf{Y})$ if $X_r \leq b_\bullet Y_r$ for some fixed deterministic constant b_\bullet , $t \geq 0$, and we write $\mathbf{X} \leq o(\mathbf{Y})$ if $\lim_{r \rightarrow \infty} X_r/Y_r = 0$ a.s. The constant b_\bullet is assumed fixed throughout.

PROPOSITION 4.2. *Let \mathbf{X} be a real-valued i.i.d. process with common mean $m = E[X_i] > 0$ and moment generating function bounded in a neighborhood of the origin.*

There exists $I_0 > \infty, \delta_0 > 0, B_0 < \infty$ such that for all $0 < \delta \leq \delta_0$, we have the following:

(i) For any $N \geq 1$, writing $S_T := \sum_{i=1}^T X_i$,

$$\mathbb{P}\left\{\inf_{T \geq 1} (S_T - (m - \delta)T) \leq -N\right\} \leq B_0 \exp(-I_0 \delta^2 N).$$

(ii) Let \mathbf{Y} be the undelayed renewal process with increment process \mathbf{X} . There exists $B_1 < \infty$ such that for $k_0 \geq 2$,

$$\lim_{r \rightarrow \infty} \sup_{0 \leq s \leq t \leq r^{k_0}} \left(\frac{Y(t) - Y(s) - (t - s)(m^{-1} + \delta)}{\log(r)} \right) \leq B_1 k_0 \delta^{-2} \quad \text{a.s.} \quad \square$$

Throughout this section we let $[\widehat{\mathbf{Q}}, \widehat{\mathbf{Z}}]$ denote any feasible trajectory for the relaxed stochastic model. It is defined on the same sample space through identical generating processes $(\mathbf{A}, \mathbf{R}, \mathbf{S})$. Our goal is to construct a policy for (2.1) that uniformly outperforms any such feasible trajectory on a time-window of the form $[\underline{T}_{r\bullet}, T_{r\bullet}]$, where

$$(4.9) \quad \underline{T}_{r\bullet} = b_0[\|x - \mathcal{P}^*(x)\| + \log(r)], \quad T_{r\bullet} = r^3,$$

with $b_0 < \infty$ sufficiently large.

The following two uniform bounds will be established for the policies constructed below, and for the optimal policy. Property P1 appears to be desirable for any network and any cost function on \mathbf{X} . However, Property P2 is desirable only when the effective cost is monotone.

Recall that $\widehat{\mathbf{w}}^*$ denotes the minimal solution to the workload relaxation (3.17), where the disturbance \mathbf{d}_0 is defined in (4.8).

Property P1 (relative optimization). For any $x \in \mathbf{X}, r \geq 1$,

$$\|\widehat{\mathbf{Q}}(t; x) - \mathcal{P}^*(\widehat{\mathbf{Q}}(t; x))\| \leq O(\log(r)) + o(1), \quad \underline{T}_{r\bullet} \leq t \leq T_{r\bullet}, \quad \text{a.s.}$$

Property P2 (relative minimal workload). For any $x \in \mathbf{X}, r \geq 1$,

$$\widehat{W}(t; x) - \widehat{w}^*(t; x) \leq O(\log(r)) + o(1), \quad \underline{T}_{r\bullet} \leq t \leq T_{r\bullet}, \quad \text{a.s.}$$

THEOREM 4.3. *Suppose that $\ell \geq \ell_u$ in (2.1) and the following additional assumptions hold:*

(i) *Assumption M holds with $n = \ell_b$, and the effective cost \bar{c} for the ℓ_b th workload-relaxation is monotone.*

(ii) *Assumptions H, S, and U hold, and the matrix B has rank ℓ_u .*

(iii) *The pair $[\mathbf{Q}^{r\circ}, \mathbf{Z}^{r\circ}]$ satisfies conditions P1 and P2.*

Then, as $r \rightarrow \infty$,

$$\sup_{[\widehat{\mathbf{Q}}^r, \widehat{\mathbf{Z}}^r]} \left(\sup_{0 \leq T \leq T_{r\bullet}} \left(\frac{1}{T} \int_0^T [c(\mathbf{Q}^{r\circ}(t; x)) - c(\widehat{\mathbf{Q}}^r(t; x))] dt \right) \right) \leq O(\log(r)) + o(1).$$

Proof. Given the allocation $\mathbf{Z}^{r\circ}$, and any other allocation $\widehat{\mathbf{Z}}^r$ satisfying (4.7), we can construct respective pseudodisturbances $\mathbf{d}_0^{r\circ}, \mathbf{d}_0^r$ via (4.8).

The proof of Theorem 4.3 is based upon a comparison of the respective optimal solutions to the ℓ_b th fluid-model relaxation (3.17), denoted $[\widehat{\mathbf{q}}^{r\circ*}, \widehat{\mathbf{z}}^{r\circ*}]$ and $[\widehat{\mathbf{q}}^{r*}, \widehat{\mathbf{z}}^{r*}]$.

This comparison is made possible via the following ‘‘coupling property’’: For any $x \in \mathsf{X}$, and all small $\delta > 0$,

$$(4.10) \quad \|d_0^{r^\circ}(t) - d_0^r(t)\| \leq \delta \|Z^{r^\circ}(t) - \widehat{Z}^r(t)\| + O(\delta^{-2} \log(r)) + o(1), \quad 0 \leq t \leq T_{r^\bullet}, \text{ a.s.}$$

This bound follows directly from Proposition 4.2 and Assumption S.

Rather than a general allocation, for each r we consider a ‘‘near-optimal’’ solution $[\widehat{Q}^r, \widehat{Z}^r]$ defined as follows. We fix $0 < T_\bullet \leq T_{r^\bullet}$, and we assume that for any other solution $[\widehat{Q}, \widehat{Z}]$,

$$\frac{1}{T_\bullet} \int_0^{T_\bullet} c(\widehat{Q}^r(t; x)) dt \leq \frac{1}{T_\bullet} \int_0^{T_\bullet} c(\widehat{Q}(t; x)) dt + O(\log(r)).$$

Recall that in this notation $O(\log(r)) \leq b_\bullet \log(r)$ with b_\bullet fixed throughout, so that the above bound is uniform in $\{\widehat{Q}\}$. It is shown in Proposition B.2 in Appendix B that a solution can be chosen so that $[\widehat{Q}^r, \widehat{Z}^r]$ satisfies conditions P1 and P2.

Combining conditions P1 and P2 for $[\widehat{Q}^{r^\circ}, \widehat{Z}^{r^\circ}]$ and $[\widehat{Q}^r, \widehat{Z}^r]$ gives

$$(4.11) \quad \begin{aligned} \|Q^{r^\circ}(t; x) - \widehat{q}^{r^\circ*}(t; x)\| &\leq O(\log(r)) + o(1), \\ \|\widehat{Q}^r(t; x) - \widehat{q}^{r*}(t; x)\| &\leq O(\log(r)) + o(1) \quad \text{a.s.} \end{aligned}$$

Theorem 3.10 and Assumption U give

$$(4.12) \quad \|\widehat{q}^{r^\circ*}(t; x) - \widehat{q}^{r*}(t; x)\| \leq O(\|d_0^{r^\circ} - d_0^r\|_\infty^t).$$

Combining (4.11), (4.12) with (4.10) and the rank condition on B then gives

$$\begin{aligned} \|Z^{r^\circ}(t; x) - \widehat{Z}^r(t; x)\| &\leq O(\|Q^{r^\circ}(t; x) - \widehat{Q}^r(t; x)\|) + O(\|d_0^{r^\circ}(t) - d_0^r(t)\|) \\ &\leq O(\|d_0^{r^\circ} - d_0^r\|_\infty^t) + O(\log(r)) + o(1) \\ &\leq \frac{1}{2} \|Z^{r^\circ} - \widehat{Z}^r\|_\infty^t + O(\log(r)) + o(1) \\ &\leq \frac{1}{2} \|Z^{r^\circ} - \widehat{Z}^r\|_\infty^{T_\bullet} + O(\log(r)) + o(1) \end{aligned}$$

uniformly for $0 \leq t \leq T_\bullet$. It follows that $\|Z^{r^\circ} - \widehat{Z}^r\|_\infty^{T_\bullet} = O(\log(r))$, and this easily implies the result. \square

The proof of Theorem 4.3 hinges on uniqueness of $[\widehat{q}^*, \widehat{z}^*]$ for a given disturbance \mathbf{d} . Without uniqueness one can attempt to search over optimal fluid allocations whose associated translation Z^{r° has minimal cost, as in the ‘‘two-armed bandit’’ (4.6). The monotonicity assumption is also critical and, as we have seen, often fails in complex network models when the workload dimension is taken larger than one. We return to this issue in section 4.5.

How then can we design a policy that satisfies conditions P1 and P2? We present here an approach based on a ‘‘discrete-review’’ structure, following [27, 33, 1]. Let $T_r > 0$, $\bar{x}^r \in \mathsf{X}$ denote, respectively, a planning horizon and safety-stock levels for the r th network. We take the explicit form

$$(4.13) \quad T_r = K_0 \log(r), \quad \bar{x}_i^r = K_1 \log(r) \bar{x}_i, \quad r \geq 1, \quad 1 \leq i \leq \ell,$$

where K_j , $j = 0, 1$, and \bar{x}_i , $i = 0, \dots, \ell$, are strictly positive constants. The ratio $\Delta_0 := K_1/K_0$ determines the likelihood of starvation.

In practice, taking a fixed safety-stock level is neither desirable nor practical—a fixed value \bar{x}^r is chosen for convenience. A more desirable choice may be a “moving target,” such as

$$\bar{x}^r = K_1 \log(\|x\| + 1) \bar{x},$$

where x is the current state of the network. It is also not necessary to assume strict positivity of *every* element of \bar{x} : it is only necessary to assume that every pooled-resource, for $i \leq \ell_r$, can work at capacity at time t if $q(t; x) \geq \bar{x}$. The results below can be extended to cover such generalizations.

We choose the allocation rates exactly as in the fluid-translation (4.5), except that we introduce safety-stock constraints that may be viewed as a shift of the origin. Let $\bar{w}^r = \widehat{\Xi} \bar{x}^r$, $r \geq 1$, and denote by $[w]_+^r$ the projection of w , in the standard ℓ_2 norm, onto the set $\widehat{W}_w^r := \{\widehat{w} + \bar{w}^r : \widehat{w} \in \widehat{W}, \widehat{w} + \bar{w}^r \geq w\}$. This is equal to the pointwise minimal element of this set under Assumption M.

Fix $\delta_0 > 0$, and consider the following linear program:

$$(4.14) \quad \begin{aligned} & \min \quad \gamma \\ & \text{subject to} \quad \gamma \geq \langle c^i, y \rangle, \quad 1 \leq i \leq \ell_c, \\ & \quad \quad \quad y = x + Bz + \alpha^r T_r, \\ & \quad \quad \quad y_i \geq (x_i + \delta_0 \bar{x}_i^r) \wedge \bar{x}_i^r, \quad 1 \leq i \leq \ell, \\ & \quad \quad \quad \widehat{\Xi} y \leq [\widehat{\Xi} x - (\mathbf{1} - \rho) T_r]_+^r, \\ & \quad \quad \quad Cz \leq T_r \mathbf{1}, \\ & \quad \quad \quad z \geq \theta. \end{aligned}$$

We assumed in Assumption M that the workload vectors satisfy $\{\xi^i : 1 \leq i \leq \ell_r\} \subset \mathbb{R}_+^\ell$. Under this condition, an application of Lemma A.1 implies that $\delta_0 > 0$ may be chosen sufficiently small so that this linear program is feasible for any $r \geq 1$.

Given a solution z^* to (4.14) we then set $\zeta^{r^\circ} := z^*/T_r$, and $Z^{r^\circ}(t; x) = t\zeta^{r^\circ}$, $0 \leq t \leq T_r$. In practice, additional constraints on Z will force an approximation, but this will be negligible for large r . This can then be repeated for each interval $[kT_r, (k+1)T_r]$ for $k \geq 0$ to obtain $(Q^{r^\circ}(t), Z^{r^\circ}(t))$ for $t \geq 0$. On any time-interval $[kT_r, (k+1)T_r]$ the buffers behave like decoupled $G/G/1$ queues.

In addition to feasibility of the linear program (4.14), the definition of Z^{r° requires *feasibility of the resulting state trajectory* so that $Q^{r^\circ}(t; x) \in X$ for all $t \geq 0$. Positivity of Q and approximate optimality follow from the large deviation bound in Proposition 4.2. We demonstrate this, and provide conditions under which P1–P2 also hold in the following two subsections.

4.4. One-dimensional workload. In this case there is a single set of pooled bottleneck-resources to be considered, and we set $\xi = \xi^1$, $\mathcal{R}^\circ = \mathcal{R}_1^\circ$. This case is special since the effective cost is always monotone, and the relaxed control problem admits a simple, pointwise optimal solution (see Proposition 3.2).

Recall that b_0 determines the times $\underline{T}_{r^\bullet}$, and $\Delta_0 = K_1/K_0$ (see (4.13)).

THEOREM 4.4. *Suppose that the following assumptions hold:*

- (i) *Assumption M holds with $n = \ell_b = 1$.*
- (ii) *Assumptions H, S, and U hold.*

Then for all $\Delta_0 > 0$ sufficiently large, there exists $b_0 < \infty$ such that Properties P1 and P2 hold for the policy defined via the linear program (4.14).

Proof. In the one-dimensional case, provided $\langle \xi, x \rangle \geq 2\bar{w}^r := 2\langle \xi, \bar{x}^r \rangle$, the linear program (4.14) to obtain $\zeta^{r\circ}$ on $[0, T_r]$ reduces to

$$(4.15) \quad \begin{aligned} \min \quad & \gamma \\ \text{subject to} \quad & \gamma \geq \langle c^i, y \rangle, \quad 1 \leq i \leq \ell_c, \\ & y = x + Bz + \alpha^r T_r, \\ & y_i \geq (x_i + \delta_0 \bar{x}_i^r) \wedge \bar{x}_i^r, \quad 1 \leq i \leq \ell, \\ & Cz \leq T_r \mathbf{1}, \\ & z \geq \boldsymbol{\theta}, \\ & \langle \xi, Bz \rangle = -T_r. \end{aligned}$$

Given a solution z^* to (4.15) we then set $\zeta^{r\circ} := T_r^{-1} z^*$. Let $y^{r\circ}$ denote the associated optimal state, starting from the initial condition x :

$$y^{r\circ} = x + (B\zeta^{r\circ} + \alpha^r)T_r.$$

As in the deterministic setting, we consider the error process

$$(4.16) \quad E^{r\circ}(t; x) := Q^{r\circ}(t; x) - \mathcal{P}^*(Q^{r\circ}(t; x)).$$

One can show as in Theorem 4.1 that for some fixed $\delta > 0$ independent of r , whenever $\|E^{r\circ}(0)\| = \|x - \mathcal{P}^*(x)\| \geq 2\|\bar{x}^r\|$, and $T^*(x) \geq T^*(\bar{x}^r) + T_r$,

$$(4.17) \quad \begin{aligned} \|y^{r\circ} - \mathcal{P}^*(y^{r\circ})\| &\leq \|E^{r\circ}(0)\| - \delta T_r, \\ \mathbb{E}[\|E^{r\circ}((k+1)T_r)\| \mid \mathcal{H}_{kT_r}] &\leq \|E^{r\circ}(kT_r)\| - \delta T_r. \end{aligned}$$

We will show that this implies P1, and that the constraint $\langle \xi, Bz \rangle = -T_r$ in (4.15) implies P2.

For $k \geq 1$ let $G_{r,k}$ denote the union of ‘‘good events,’’

$$\begin{aligned} G_{r,k} = & \left\{ W^{r\circ}(kT_r) \leq 2\bar{w}^r \right\} \\ & \bigcup \left\{ \frac{1}{2}\delta_0 \bar{x}^r \leq E^{r\circ}(t; x) \leq \frac{3}{2}\bar{x}^r, \quad kT_r \leq t \leq (k+1)T_r, \right. \\ & \left. \text{and } \langle \xi, B[Z^{r\circ}((k+1)T_r) - Z^{r\circ}(kT_r)] \rangle = -T_r \right\}, \end{aligned}$$

and define for any r

$$G_r = \bigcup_{\underline{T}_{r\bullet} \leq k \leq \underline{T}_{r\bullet}/T_r} G_{r,k}.$$

For sufficiently large b_0 and constants $B_2 < \infty$, $I_2 > 0$, Proposition 4.2 implies the bound $\mathbb{P}\{G_r^c\} \leq r^3 B_2 \exp(-I_2 \Delta_0 \log(r))$, $r \geq 1$. For $\Delta_0 \geq 5I_2^{-1}$ this is bounded by $B_2 r^{-2}$, and it then follows that

$$\sum_{r=1}^{\infty} \mathbb{P}\{G_r^c\} < \infty.$$

From the Borel–Cantelli lemma we can conclude that, with probability one, each state-allocation trajectory $[Q^{r\circ}, Z^{r\circ}]$ eventually satisfies G_r for large enough r . It follows that P1 and P2 also hold and that $Q^{r\circ}$ evolves in \mathbb{X} for all large r . \square

4.5. Higher dimensions. For the general workload dimension, even if the fluid model admits a pointwise optimal solution, one cannot hope to obtain the strong sample-path optimality established in Theorem 4.4 for the stochastic model. Consider the workload processes

$$W^r(t; x) = \widehat{\Xi}Q^r(t; x), \quad W^{r^\circ}(t; x) = \widehat{\Xi}Q^{r^\circ}(t; x).$$

In heavy traffic, any greedy policy would attempt to drive $W^r(t; x)$ into the set \widehat{W}^+ . This is illustrated in Figure 10. Initially, the trajectory W^{r° mimics the behavior of the fluid model. It is probable that a sample path will then drift throughout the region \widehat{W}^+ if $\rho \sim 1$. For the sample path shown, initially $\bar{c}(W^r(t; x))$ is much greater than $\bar{c}(W^{r^\circ}(t; x))$, but then the opposite is true following the upward drift of $W_2^{r^\circ}$ shown in the figure.

This counterexample depends upon the specific geometry shown. Although a pointwise optimal solution exists for the fluid model, the effective cost \bar{c} is not monotone. Consequently, property P2 is not desirable—the optimal workload trajectory \widehat{w}^* for the fluid model is not pointwise minimal.

Assuming that the effective cost is monotone, the arguments used in the proofs of Theorem 3.10 and Theorem 4.4 may be applied to establish the following consequences.

THEOREM 4.5. *Suppose that the following assumptions hold:*

- (i) *Assumption M holds with $n = \ell_b$, and the effective cost \bar{c} for the ℓ_b th workload-relaxation is monotone.*
- (ii) *Assumptions H, S, and U hold.*

Then for all $\Delta_0 > 0$ sufficiently large, there exists $b_0 < \infty$ such that Properties P1 and P2 hold for the policy defined via the linear program (4.14).

Proof. The proof of condition P1, and positivity of the state trajectory Q^{r° , is identical to the proof in the one-dimensional case since (4.17) continues to hold for the associated error process given in (4.16).

To establish condition P2 we first note that the allocation rate $\zeta_k^{r^\circ}$ on $[kT_r, (k+1)T_r]$ is designed to be *almost optimal* for a disturbance-free model on $[kT_r, (k+1)T_r]$,




FIG. 10. The figure shows two trajectories for the stochastic workloads $W^r(t)$ and $W^{r^\circ}(t)$, $t \geq 0$.

in the sense that

$$(4.18) \quad W^{r\circ}(kT_r; x) + \widehat{\Xi}[B\zeta_k^{r\circ}T_r + \alpha^r T_r] \leq [W^{r\circ}(kT_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r]_+^r.$$

Using (4.18) we verify the following restricted form of condition P2 by induction:

$$(4.19) \quad W^{r\circ}(kT_r; x) \leq \widehat{w}^{r\circ*}(kT_r; x) + 2\bar{w}^r, \quad k \geq 0.$$

In fact, this bound combined with Proposition 4.2(ii) implies that P2 holds.

The inequality (4.19) is obvious for $k = 0$. Assuming it is valid for a given $k \geq 1$, denote $D(k+1) := d^{r\circ}((k+1)T_r) - d^{r\circ}(kT_r)$, and consider the following bounds:

$$\begin{aligned} W^{r\circ}((k+1)T_r; x) &\leq [W^{r\circ}(kT_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r]_+^r + D(k+1) \\ &\quad \text{by (4.18)} \\ &\leq [\widehat{w}^{r\circ*}(kT_r; x) + 2\bar{w}^r - (\mathbf{1} - \boldsymbol{\rho}^r)T_r]_+^r + D(k+1) \\ &\quad \text{by induction} \\ &= \widehat{w}^{r\circ*}(kT_r; x) + 2\bar{w}^r - (\mathbf{1} - \boldsymbol{\rho}^r)T_r + D(k+1) \\ &\quad \text{for } r \text{ sufficiently large} \\ &\leq [\widehat{w}^{r\circ*}(kT_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r + D(k+1)]_+ + 2\bar{w}^r \\ &\leq \widehat{w}^{r\circ*}((k+1)T_r; x) + 2\bar{w}^r \\ &\quad \text{by Lemma 3.9.} \end{aligned}$$

The equality in the third line follows since \bar{w}^r is an interior point of \widehat{W} , and $\boldsymbol{\rho}^r \rightarrow \mathbf{1}$ as $r \rightarrow \infty$. This completes the verification of the induction hypothesis. \square

5. Conclusions. The results of this paper lead to practical control solutions for large networks. One must consider an appropriate relaxation for the fluid model, define idealized target states through this idealization, and use safety stocks and some regulation policy to attempt to meet these target values.

Consider for example the network illustrated in Figure 1. For certain parameters a two-dimensional workload-relaxation is justifiable, and a policy that is nearly optimal in heavy traffic can be computed by hand once the effective cost is found. When the cost c is linear and $\mathbf{X} = \mathbb{R}_+^\ell$, the effective cost $\bar{c}(x)$ is the value of the linear program

$$(5.1) \quad \begin{aligned} \min \quad &\langle c, x \rangle \\ \text{subject to} \quad &\xi_i^T x = w_i, \quad i = 1, 2, \\ &x \geq \boldsymbol{\theta}. \end{aligned}$$

It is amazing that optimal policy synthesis can be conceptualized so easily for such a complex model!

In practice, it may be difficult to summarize *every* goal in a single cost function. Optimization may be viewed as a method of generating a family of candidate *good policies*. One can then choose among these or formulate generalizations to satisfy various goals.

A complex network model such as that shown in Figure 1 illustrates the importance of taking a flexible viewpoint in modelling, and in control design. By restricting to a basic feasible solution of (5.1), one may assume that an optimal trajectory $\widehat{q}^*(t; x)$ is null, with the exception of at most two positive components when the cost function c is linear. This is born out in numerical experiments conducted in [14]. After a short

transient period, it is observed that in all but two of the buffers, all of the inventory vanishes in the optimal fluid state trajectory. Similar behavior is commonly seen in the heavy-traffic networks literature (see, e.g., [29, 25, 1]).

In practice, such behavior is rarely feasible because buffers are finite. One can add a state space constraint to both models: for an $\ell_s \times \ell$ constraint matrix C_s ,

$$(5.2) \quad C_s Q(t; x) \leq \mathbf{1}, \quad Q(t; x) \geq \boldsymbol{\theta}, \quad t \geq 0.$$

The state space is then redefined via $X = \{x \in \mathbb{R}^\ell : x \geq \boldsymbol{\theta}, C_s x \leq \mathbf{1}\}$, and in this case the set \widehat{W} given in (3.6) is no longer a simple positive cone. These additional constraints increase the complexity of optimal state trajectories so that work is distributed across a greater number of buffers. Alternatively, if a strictly convex cost function is used, rather than a linear one, then more reasonable optimal trajectories will be obtained.

Another question concerns uncertainty. In telecommunications applications one may know little about the arrival rates to the system, and in a manufacturing application *demand* may be uncertain. One approach is to define a set of *generalized Klimov indices*, as in Proposition 12 of [36]. Alternatively, given prior information regarding arrival rates, one can expand the definition of V . Suppose that A is a polyhedron that defines possible arrival rates. The corresponding worst-case emptying time is given by

$$\bar{T}^*(x) = \max_{\alpha \in A} \max_{1 \leq i \leq \ell_r} \frac{\langle \xi^i, x \rangle}{1 - \langle \xi^i, \alpha \rangle}.$$

It is then straightforward to design efficient policies for the fluid model that drain the system before this time without knowledge of the exact value of α . Other approaches have been considered recently in [31, 40, 19].

It has been taken for granted in this paper that activities and buffers far outnumber resources. However, in communication applications, particularly in wireless models, one frequently finds that the set of possible allocation rates is a highly complex convex set (see, e.g., [41, 44]). In particular, it may not be a polyhedron. One interpretation is that in wireless models there are an infinite number of resources through time-division, frequency selection, multiple paths, or choices in coding schemes. Extensions of the methods developed here may be possible provided the rate set V is suitably smooth, and in this case a one-dimensional relaxation is suggested.

There are many questions left unanswered.

(i) Can one formulate a version of Theorem 4.5 when the fluid model admits pointwise optimal solutions, yet the effective cost \bar{c} is not monotone? This question is interesting even in the case of a single bottleneck since sample-path optimality does not hold if ξ^1 has any negative components (see [7]).

(ii) What if a pointwise optimal allocation does not exist for the ℓ_i th workload-relaxation? Can one obtain a near-optimal policy in this case (in the infinite-horizon sense (2.3))?

(iii) The policies described in this paper are based on state-feedback, using a workload-based model. It is expected that RBM models will play a role in the determination of optimality in the mean and in a finer performance analysis.

(iv) Can efficient recursive algorithms, based on workload dimension, be constructed for policy synthesis on a fluid scale?

(v) Where are the sources of highest sensitivity in control design?

(vi) Do the results of this paper lead to improved methods for performance approximation via simulation, or through calculation, by exploiting the simplicity of the network model following state space collapse?

(vii) Finally, extensions of these algorithms have been formulated for sequencing and routing in the face of breakdowns or preventative maintenance. We are eager to see how these methods actually work in practice.

Topics (i)–(v) are considered in what follows [7, 28], but the story is far from complete.

Appendix A. Workload relaxations. The proof of Theorem 4.1 is based on the following lemma, which shows that, relative to the system load, exchangeable states for the ℓ_b th workload-relaxation are almost exchangeable for the original fluid model when r is large.

LEMMA A.1. *Suppose that Assumption H holds. There exists $b_0 < \infty$ such that for any $x, y \in \mathsf{X}$, and any $r \geq 1$, the time to reach y from x is uniformly bounded,*

$$T^*(x, y) \leq b_0 \|x - y\| \quad \text{whenever} \quad \widehat{\Xi}(y - x) \geq \boldsymbol{\theta}.$$

Proof. If $\langle \xi^i, y - x \rangle \geq 0$ for $1 \leq i \leq \ell_b$, then it follows from the definition of T^* that

$$\begin{aligned} T^*(x, y) &= \max_{i \geq 1} \frac{\langle \xi^i, x - y \rangle}{1 - \langle \xi^i, \alpha^r \rangle} \\ &= \max_{i > \ell_b} \frac{\langle \xi^i, x - y \rangle}{1 - \langle \xi^i, \alpha^r \rangle}, \quad 1 \leq r < \infty. \end{aligned}$$

The right-hand side is bounded in r by construction of α^r and the definition of ℓ_b . It also scales linearly on scaling $(x - y)$. This gives the required bound. \square

Proof of Theorem 4.1. Fix $0 < t < \underline{T}_{r_0}(x)$, and define

$$\widehat{v}^\perp = -\beta \frac{e^r(t; x)}{\|e^r(t; x)\|}.$$

The constant $\beta > 0$ is chosen so that \widehat{v}^\perp is a boundary point of V_r . We have the explicit formula $\beta^{-1} = \frac{T^*(x^1, x^2)}{\|x^2 - x^1\|}$, with $x^1 = q(t; x)$, $x^2 = \widehat{q}^*(t; x)$.

We have already remarked that the constraints ensure that (i) holds so that $w_i(t; x) \leq \widehat{w}_i^*(t; x)$ for all $t \geq 0$, and $1 \leq i \leq \ell_b$. It follows that $\widehat{\Xi} \widehat{v}^\perp \geq \boldsymbol{\theta}$, and Lemma A.1 implies directly that $\beta = \beta(r)$ is uniformly bounded from below in r . Applying this and Assumption U(i), we conclude that there is some fixed $\Delta > 0$, independent of $x \in \mathsf{X}$ and $r \geq 1$, such that for all $0 \leq t < \underline{T}_{r_0}(x)$ and sufficiently small $s > 0$,

$$(A.1) \quad c(q(t; x) + s\widehat{v}^\perp) - c(q(t; x)) \leq -\Delta s.$$

Now let

$$(A.2) \quad v = \widehat{v}^* + \left(1 - \frac{1}{2} \frac{r_0}{r}\right) \widehat{v}^\perp,$$

where $\widehat{v}^* = \frac{d}{dt} \widehat{q}^*(t; x)$ and r_0 is a constant. We show that this is in V_r for any $r \geq r_0$ provided r_0 is sufficiently large. For $1 \leq i \leq \ell_b$ we have $\langle \xi^i, \widehat{v}^\perp \rangle \geq 0$, and hence for $r \geq r_0$,

$$\langle \xi^i, v \rangle := \langle \xi^i, \widehat{v}^* + (1 - r_0/(2r)) \widehat{v}^\perp \rangle \geq \langle \xi^i, \widehat{v}^* \rangle \geq -(1 - \rho_i^r).$$

For $i > \ell_b$ we can reason as follows: The identity (4.4) implies that $\|\frac{d}{dt}\widehat{w}^*(t)\| \leq K_0/r$ for some $K_0 < \infty$ and all $t > 0$. Since \mathcal{X}^* is continuous, we must have a similar bound for \widehat{q}^* , so that $\|\widehat{v}^*\| \leq K_1/r$ for some finite K_1 . Then, for $i > \ell_b$ and $r \geq r_0$,

$$\begin{aligned} \langle \xi^i, v \rangle &\geq \langle \xi^i, \widehat{v}^* \rangle + \frac{1}{2} \frac{r_0}{r} (1 - \rho_i^r) - (1 - \rho_i^r) \\ &\geq \frac{1}{r} (\frac{1}{2} r_0 (1 - \rho_i^r) - K_1 \|\xi^i\|) - (1 - \rho_i^r). \end{aligned}$$

Hence, to ensure feasibility of v , it is enough to choose $r_0 > 2K_1 \|\xi^i\| (1 - \rho_i^\infty)^{-1}$ for all $i > \ell_b$.

Using (A.1), (A.2), and the minimality of $\langle \nabla c, \frac{d}{dt}q \rangle$, we obtain the bound, for any $0 < t < \underline{T}_{r_0}$, $r \geq r_0$,

$$\frac{d}{dt}(c(q(t; x)) - c(\widehat{q}^*(t; x))) \leq -\left(1 - \frac{1}{2} \frac{r_0}{r}\right) \Delta \leq -\frac{1}{2} \Delta.$$

Let $g(t) = c(q(t; x)) - c(\widehat{q}^*(t; x))$, $t \geq 0$. This function is piecewise linear on $(0, \infty)$ and satisfies $g(0+) = c(x) - c(\mathcal{P}^*(x))$, where \mathcal{P}^* is defined in (3.10). The previous bound then gives

$$g(t) \leq g(0+) - \frac{1}{2} \Delta t, \quad 0 < t < \underline{T}_{r_0},$$

and since g is nonnegative, $g(t) = 0$ for $t \geq 2g(0+)/\Delta$. Assumption U(ii) then implies that $q(t; x) = \widehat{q}^*(t; x)$ for such t , and hence

$$\underline{T}_{r_0}(x) < \frac{2}{\Delta} g(0+) = \frac{2}{\Delta} (c(x) - c(\mathcal{P}^*(x))).$$

This proves (ii) and (iii) since c is a norm, and results (iv) and (v) follow immediately. \square

Appendix B. Stochastic models.

Proof of Proposition 4.2. In part (i) we are asking, *When can the graph (T, S_T) of the partial sums of X_i hit the line $l(T) = (m - \delta)T - N$ for some $T \geq 0$?* Hence, the bound in (i) follows easily from Cramer's theorem [12].

For (ii) we define, for any $i \geq 1$, the event

$$\mathcal{E}_i := \{Y(S_i + T) - Y(S_i) - m^{-1}T \geq \delta T + N \text{ some } T \geq 0\}.$$

Using the fact that $S_i := \sum_{j=1}^i X_j$, $i \geq 1$, is equal to the time of the i th jump of \mathbf{Y} , we obtain the identity

$$\mathcal{E}_i = \left\{ \sum_{j=i+1}^{(m^{-1}+\delta)T+N} X_j \leq T, \text{ some } T \geq 0 \right\}.$$

Applying (i) gives a bound of the form $\mathbb{P}(\mathcal{E}_i) \leq B_1 \exp(-I_1 \delta^2 N)$ for some $B_1 < \infty$, $I_1 > 0$, and all i, N . Consequently, for any $r \geq 1$,

$$\begin{aligned} &\mathbb{P} \left\{ \sup_{0 \leq s \leq t \leq r^{k_0}} \left(\frac{Y(t) - Y(s) - (t-s)(m^{-1} + \delta)}{N} \right) \geq 1 \right\} \\ &\leq \mathbb{P}\{S_{r^{k_0+1}} \leq r^{k_0}\} + \sum_{i=1}^{r^{k_0+1}} \mathbb{P}\{\mathcal{E}_i\} \\ &\leq B_1 \exp(-I_1 \delta_0 r) + r^{k_0+1} B_1 \exp(-I_1 \delta^2 N). \end{aligned}$$

We now define N via $I_1 \delta^2 N = \log(r^{k_0+3})$, so that the right-hand side is bounded by $2B_2 r^{-2}$. This is summable, and hence by the Borel–Cantelli lemma,

$$\limsup_{r \rightarrow \infty} \left\{ \sup_{0 \leq s \leq t \leq r^{k_0}} \left(\frac{Y(t) - Y(s) - (t-s)(m^{-1} + \delta)}{(I_1 \delta^2)^{-2} \log(r^{k_0+3})} \right) \right\} \leq 1 \quad \text{a.s.} \quad \square$$

We may now develop properties of the stochastic relaxation used to prove Theorem 4.3. Consider the value function for the relaxed stochastic network for a given $T > 0$, $x \in \mathbf{X}$,

$$\widehat{\Gamma}^{r*}(T; x) = \min \frac{1}{T} \int_0^T c(\widehat{Q}(t)) dt,$$

where the minimum is over all allocations \widehat{Z} satisfying the constraints (4.7), subject to the additional constraint that \widehat{Q} evolves in $\mathbf{X} = \mathbb{R}_+^\ell$.

Given the value of \widehat{Z}^* at time t , the associated idleness process is given by

$$\widehat{\Gamma}^{r*}(t; x) = t\mathbf{1} - \widehat{C}\widehat{Z}^{r*}(t; x), \quad t \geq 0, x \in \mathbf{X}.$$

Conversely, one can determine the optimal allocation given the current value of the idleness. If $\widehat{\Gamma}^{r*}(t; x) = \widehat{I}$, then we take $\widehat{Z}^{r*}(t; x)$ to be the minimizer \widehat{Z}^{r*} of the nonlinear optimization problem

$$(B.1) \quad \begin{array}{ll} \min c(y) & \text{subject to} \\ & y = x - S(\widehat{Z}) + R(\widehat{Z}) + A(t), \\ & \widehat{C}\widehat{Z} = \widehat{I}, \\ & \widehat{Z} \in \mathbb{R}^{\ell_u}, \\ & y \geq \boldsymbol{\theta}. \end{array}$$

This representation leads to the following conclusion.

PROPOSITION B.1. *Suppose that Assumptions H and S hold. Then, the optimal solution $[\widehat{Q}^{r*}, \widehat{Z}^{r*}]$ satisfies condition P1.*

Proof. Fix any $t_0 > 0$, set $y = \widehat{Q}^{r*}(t_0)$, let $y^* = \mathcal{P}^*(y)$, and set ζ^+ as any solution to $B\zeta^+ = y^* - y$, so that $\widehat{C}\zeta^+ = \boldsymbol{\theta}$. Under Assumption U(i) there exists $\kappa > 0$ such that $c(y + sB\zeta^+) \leq c(y) - s\|\zeta^+\|\kappa$ for $s \leq 1$.

Consider the allocation $\widehat{Z}^{*\Delta}$ given by $\widehat{Z}^{*\Delta}(t) = \widehat{Z}^{r*}(t)$ if $t \neq t_0$, and $\widehat{Z}^{r\Delta}(t_0) = \widehat{Z}^{r*}(t_0) + \Delta\zeta^+$. On the remainder of \mathbb{R}_+ we again suppose that this allocation is linear on each time-horizon. This is feasible for a range of $\Delta \geq 0$, and by Proposition 4.2 we have

$$c(\widehat{Q}^{r*}(t_0)) \geq c(\widehat{Q}^{r\Delta}(t_0)) - O(s\|\zeta^+\|\kappa) + O(\log(r)).$$

We must therefore have $\|\zeta^+\| = O(\log(r))$, so that P1 holds. \square

We may also establish a form of P2.

PROPOSITION B.2. *Suppose that Assumptions H, M, U, and S hold, where all bounds are with respect to the ℓ_b th workload-relaxation. Then for any $r \geq 1$ and any $0 < T_\bullet \leq T_{\bullet\bullet}$, there exists a solution $[\widehat{Q}^r, \widehat{Z}^r]$ that satisfies conditions P1 and P2, and*

$$(B.2) \quad \frac{1}{T_\bullet} \int_0^{T_\bullet} c(\widehat{Q}^r(t)) dt \leq \widehat{\Gamma}^{r*}(T_\bullet; x) + O(\log(r)).$$

Proof. The proof is again by comparison. We approximately retain the convention that $\widehat{\mathbf{Z}}^r$ is determined from its idleness process through the following restricted form of (B.1): For a given value $\widehat{I} = \widehat{I}^r(t; x)$ we take $\widehat{Z}^r(t; x)$ to be the minimizer \widehat{Z}^* of

$$(B.3) \quad \begin{array}{ll} \min c(y) & \text{subject to} \\ & y = x - S(\widehat{Z}) + R(\widehat{Z}) + A(t), \\ & \widehat{C}\widehat{Z} = \widehat{I}, \\ & \widehat{Z} \in \mathbb{R}^{\ell_u}, \\ & y \geq \frac{1}{2}\bar{x}^r. \end{array}$$

Under this restriction, one can adapt the proof given in Proposition B.1 to show that $[\widehat{\mathbf{Q}}^r, \widehat{\mathbf{Z}}^r]$ satisfies Property P1.

We now show how $\widehat{\mathbf{I}}^r$ may be constructed so that P2 holds.

For a given $\Delta > 0$, let $\widehat{\mathbf{I}}^{r\Delta}$ denote the idleness process defined by $\widehat{I}^{r\Delta}(t; x) = \widehat{I}^{r*}(t; x) + \Delta\bar{w}^r$, $t \geq 0$, where $\bar{w}^r := \widehat{\Xi}\bar{x}^r$. An application of Proposition 4.2 implies that Δ can be chosen so large that the resulting state trajectory satisfies

$$\widehat{W}^{r*}(t; x) + 2\Delta\bar{w}^r + o(1) \geq \widehat{W}^{r\Delta}(t; x) \geq \bar{w}^r - o(1), \quad 0 \leq t \leq T_{\bullet}.$$

This can be chosen independently of $r \geq 1$ and independently of $0 < T_{\bullet} \leq T_{r\bullet}$. It is obvious that (B.2) holds for the allocations $\{\widehat{\mathbf{I}}^{r\Delta}\}$.

We now refine this allocation to form an allocation $\widehat{\mathbf{I}}^r$ as follows. We first define this for $t = kT_r$ and then take it to be linear on each interval $[kT_r, (k+1)T_r]$ for each $k \geq 0$.

For $k = 0$ we set $\widehat{I}^r(0; x) := \widehat{I}^{r\Delta}(0; x) = \Delta\bar{w}^r$. For all $k \geq 1$, given that $\widehat{I}^r((k-1)T_r; x)$ has been specified, we define an upperbound $\bar{I} \in \mathbb{R}_+^{\ell_b}$ on the idleness rate on the interval $[(k-1)T_r, kT_r]$. This is given as the solution to

$$(B.4) \quad \widehat{W}^r((k-1)T_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r + \bar{I}T_r = [\widehat{W}^r((k-1)T_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r]_+^r.$$

We then define $\widehat{I}^r(kT_r; x)$ as

$$\widehat{I}^r(kT_r; x) := \min\left([\widehat{I}^r((k-1)T_r; x) + \bar{I}T_r], \widehat{I}^{r\Delta}(kT_r; x)\right),$$

where the minimum is interpreted componentwise. One can show that for sufficiently large $r \geq 1$, $\widehat{W}^r(t; x) - \frac{1}{2}\bar{w}^r \in \widehat{W}$, $t \geq 0$. It follows that $\widehat{\mathbf{I}}^r$ defines a feasible allocation \mathbf{Z}^r , in the sense that the nonlinear program (B.3) is feasible when $\widehat{I} = \widehat{I}^r(t)$. Following familiar arguments we may conclude that the resulting state trajectory $\widehat{\mathbf{Q}}^r$ evolves in \mathbf{X} for sufficiently large r .

Moreover, (B.4) implies that the resulting workload process satisfies a bound similar to (4.18): We have by construction

$$\widehat{W}^r(kT_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r + \hat{i}_k^r T_r \leq [W^r(kT_r; x) - (\mathbf{1} - \boldsymbol{\rho}^r)T_r]_+^r,$$

where $\hat{i}_k^r := T_r^{-1}[\widehat{I}^r((k+1)T_r; x) - \widehat{I}^r(kT_r; x)]$ denotes the idleness on the k th interval. The proof of P2 is then identical to the proof of P2 for $\widehat{\mathbf{Q}}^{r_0}$ given in Theorem 4.5.

Finally, since $\widehat{I}^r(kT_r; x) \leq \widehat{I}^{r\Delta}(kT_r; x)$ for all k and $\widehat{\mathbf{I}}^{r\Delta}$ satisfies (B.2) by construction, we may establish (B.2) for $\widehat{\mathbf{I}}^r$ by an application of Proposition 4.2 as in the proof of Proposition B.1. \square

Acknowledgments. The author would like to express his thanks to Michael Chen, Michael Harrison, and Shane Henderson for many useful comments on an earlier

draft of this manuscript. The anonymous referee also provided valuable input. Of course, the author takes full responsibility for any remaining errors. Thanks are also due to Maury Bramson, Kavita Ramanan, and Ruth Williams for sharing their unpublished work, and for many fruitful discussions.

REFERENCES

- [1] S. L. BELL AND R. J. WILLIAMS, *Dynamic scheduling of a system with two parallel servers: Asymptotic policy in heavy traffic*, in Proceedings of the 38th IEEE Conference on Decision and Control, Vol. 3, IEEE Press, Piscataway, NJ, 1999, pp. 2255–2260.
- [2] A. BERNARD AND A. EL KHARROUBI, *Régulations déterministes et stochastiques dans le premier “orthant” de R^n* , Stochastics Stochastics Rep., 34 (1991), pp. 149–167.
- [3] D. BERTSEKAS AND R. GALLAGER, *Data Networks*, Prentice–Hall, Englewood Cliffs, NJ, 1992.
- [4] M. BRAMSON, *State space collapse with application to heavy traffic limits for multiclass queueing networks*, Queueing Systems Theory Appl., 30 (1998), pp. 89–148.
- [5] C. CHEN, Z. JIA, AND P. VARAIYA, *Causes and cures of highway congestion*, IEEE Control Systems Magazine, 21 (2001), pp. 26–32.
- [6] H. CHEN AND D. D. YAO, *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*, Appl. Math. 46, Springer–Verlag, New York, 2001.
- [7] M. CHEN, C. PANDIT, AND S. P. MEYN, *In search of sensitivity in network optimization*, Queueing Systems Theory Appl., to appear.
- [8] J. G. DAI, *On positive Harris recurrence of multiclass queueing networks: A unified approach via fluid limit models*, Ann. Appl. Probab., 5 (1995), pp. 49–77.
- [9] J. G. DAI, *A fluid-limit model criterion for instability of multiclass queueing networks*, Ann. Appl. Probab., 6 (1996), pp. 751–757.
- [10] J. G. DAI AND S. P. MEYN, *Stability and convergence of moments for multiclass queueing networks via fluid limit models*, IEEE Trans. Automat. Control, 40 (1995), pp. 1889–1904.
- [11] J. G. DAI AND G. WEISS, *A fluid heuristic for minimizing makespan in job-shops*, Oper. Res., 50 (2002), pp. 692–707.
- [12] A. DEMBO AND O. ZEITOUNI, *Large Deviations Techniques and Applications*, 2nd ed., Springer–Verlag, New York, 1998.
- [13] B. T. DOSHI, *Optimal control of the service rate in an $M/G/1$ queueing system*, Adv. Appl. Probab., 10 (1978), pp. 682–701.
- [14] M. CHEN, R. DUBRAWSKI, AND S. P. MEYN, *Management of demand-driven production systems*, IEEE Trans. Automat. Control., submitted.
- [15] P. DUPUIS AND K. RAMANAN, *Convex duality and the Skorokhod problem. I*, Probab. Theory Related Fields, 115 (1999), pp. 153–195.
- [16] P. DUPUIS AND K. RAMANAN, *Convex duality and the Skorokhod problem. II*, Probab. Theory Related Fields, 115 (1999), pp. 197–236.
- [17] P. DUPUIS AND K. RAMANAN, *An explicit formula for the solution of certain optimal control problems on domains with corners*, Teor. ĭmovir. Mat. Stat., 63 (2000), pp. 32–48.
- [18] P. DUPUIS AND K. RAMANAN, *A multiclass feedback queueing network with a regular Skorokhod problem*, Queueing System Theory Appl., 36 (2000), pp. 327–349.
- [19] A. ERYILMAZ, R. SRIKANT, AND J. R. PERKINS, *Stable scheduling policies for broadcast channels*, in Proceedings of the 2002 IEEE Symposium on Information Theory, IEEE Press, Piscataway, NJ, 2002, p. 382.
- [20] S. B. GERSHWIN, *Manufacturing Systems Engineering*, Prentice–Hall, Englewood Cliffs, NJ, 1993.
- [21] ĭ. ĭ. GĭHMAN AND A. V. SKOROHOD, *Stochastic Differential Equations*, Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 72, Kenneth Wickwire, trans., Springer–Verlag, New York, 1972.
- [22] J. M. HARRISON, *Brownian models of queueing networks with heterogeneous customer populations*, in Stochastic Differential Systems, Stochastic Control Theory and Applications (Minneapolis, Minn., 1986), Springer–Verlag, New York, 1988, pp. 147–186.
- [23] J. M. HARRISON, *Brownian models of open processing networks: Canonical representations of workload*, Ann. Appl. Probab., 10 (2000), pp. 75–103.
- [24] J. M. HARRISON, *Stochastic networks and activity analysis*, in Analytic Methods in Applied Probability: In Memory of Fridrikh Karpelevich, Amer. Math. Soc. Transl. Ser. 2 207, Yu. M. Suhov, ed., AMS, Providence, RI, 2002.
- [25] J. M. HARRISON AND J. A. VAN MIEGHEM, *Dynamic control of Brownian networks: State space*

- collapse and equivalent workload formulations*, Ann. Appl. Probab., 7 (1997), pp. 747–771.
- [26] J. M. HARRISON AND R. J. WILLIAMS, *Brownian models of open queueing networks with homogeneous customer populations*, Stochastics, 22 (1987), pp. 77–115.
- [27] J. M. HARRISON, *The BIGSTEP approach to flow management in stochastic processing networks*, in Stochastic Networks Theory and Applications, F. P. Kelly, S. Zachary, and I. Ziedins, eds., Clarendon Press, Oxford, UK, 1996, pp. 57–89.
- [28] S. G. HENDERSON, S. P. MEYN, AND V. TADIC, *Performance evaluation and policy selection in multiclass networks*, Discrete Event Dyn. Syst., 13 (2003), pp. 149–189.
- [29] F. C. KELLY AND C. N. LAWS, *Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling*, Queueing Systems Theory Appl., 13 (1993), pp. 47–86.
- [30] H. J. KUSHNER, *Heavy Traffic Analysis of Controlled Queueing and Communication Networks*, Springer-Verlag, New York, 2001.
- [31] R. LEELAHAKRIENGKRAI AND R. AGRAWAL, *Scheduling in multimedia wireless networks*, in Proceedings of the 17th International Teletraffic Congress, Brazil, 2001.
- [32] X. LUO AND D. BERTSIMAS, *A new algorithm for state-constrained separated continuous linear programs*, SIAM J. Control Optim., 37 (1998), pp. 177–210.
- [33] C. MAGLARAS, *Design of dynamic control policies for stochastic processing networks via fluid models*, in Proceedings of the 36th IEEE Conference on Decision and Control, IEEE Press, Piscataway, NJ, 1997, pp. 1208–1213.
- [34] S. P. MEYN, *Transience of multiclass queueing networks via fluid limit models*, Ann. Appl. Probab., 5 (1995), pp. 946–957.
- [35] S. P. MEYN, *Stability and optimization of queueing networks and their fluid models*, in Mathematics of Stochastic Manufacturing Systems (Williamsburg, VA, 1996), American Mathematical Society, Providence, RI, 1997, pp. 175–199.
- [36] S. P. MEYN, *Sequencing and routing in multiclass queueing networks part I: Feedback regulation*, SIAM J. Control Optim., 40 (2001), pp. 741–776.
- [37] M. C. PULLAN, *Forms of optimal solutions for separated continuous linear programs*, SIAM J. Control Optim., 33 (1995), pp. 1952–1977.
- [38] M. I. REIMAN, *Open queueing networks in heavy traffic*, Math. Oper. Res., 9 (1984), pp. 441–458.
- [39] M. I. REIMAN, *A multiclass queue in heavy traffic*, Adv. Appl. Probab., 20 (1988), pp. 179–207.
- [40] S. SHAKKOTTAI AND A. STOLYAR, *Scheduling for multiple flows sharing a time-varying channel: The exponential rule*, in Analytic Methods in Applied Probability: In Memory of Fridrikh Karpelevich, Amer. Math. Soc. Transl. Ser. 2 207, Yu. M. Suhov, ed., AMS, Providence, RI, 2002.
- [41] D. N. C. TSE AND S. V. HANLY, *Multiaccess fading channels. I. Polymatroid structure, optimal resource allocation and throughput capacities*, IEEE Trans. Inform. Theory, 44 (1998), pp. 2796–2815.
- [42] G. WEISS, *Optimal draining of fluid re-entrant lines*, in Stochastic Networks: Theory and Applications, Roy. Statist. Soc. Lecture Note Ser. 4, F. P. Kelly, S. Zachary, and I. Ziedins, eds., Oxford University Press, Oxford, 1996, pp. 19–34.
- [43] R. J. WILLIAMS, *Diffusion approximations for open multiclass queueing networks: Sufficient conditions involving state space collapse*, Queueing Systems Theory Appl., 30 (1998), pp. 27–88.
- [44] L. XIAO, M. JOHANSSON, H. HINDI, S. BOYD, AND A. GOLDSMITH, *Joint optimization of communication rates and linear systems*, in Proceedings of the 40th IEEE Conference on Decision and Control, Vol. 3, IEEE Press, Piscataway, NJ, 2001, pp. 2321–2326.
- [45] L. ZUCKERMAN AND M. L. WALD, *Gridlock in the skies: A special report; Crisis for air traffic system: More passengers, more delays*, New York Times, 5 September, 2000, sec. A, p. 1, late edition.