

ASYMPTOTICALLY OPTIMAL CONTROLS FOR TIME-INHOMOGENEOUS NETWORKS*

MILICA ČUDINA[†] AND KAVITA RAMANAN[‡]

Abstract. A framework is introduced for the identification of controls for single-class time-varying queueing networks that are asymptotically optimal in the so-called uniform acceleration regime. A related, but simpler, first-order (or fluid) control problem is first formulated. For a class of performance measures that satisfy a certain continuity property, it is then shown that any sequence of policies whose performances approach the infimum in the fluid control problem is asymptotically optimal for the original network problem. Examples of performance measures with this property are described, and simulations implementing asymptotically optimal policies are presented. The use of directional derivatives of the reflection map for solving fluid control problems is also illustrated. This work serves to complement a large body of literature on asymptotically optimal controls for time-homogeneous networks.

Key words. asymptotic optimality, directional derivatives, fluid limit, inhomogeneous networks, Poisson point processes, queueing networks, reflection map, stochastic optimal control, uniform acceleration

AMS subject classifications. 60K25, 60M20, 93E20

DOI. 10.1137/090762026

1. Introduction. Most real-world queueing systems evolve according to laws that vary with time. The expository paper [32] outlines the applications of time-varying stochastic networks to telecommunications. In the context of computer engineering, these networks arise in the fields of power aware scheduling and temperature aware scheduling (see, e.g., [3, 44, 45]) and the design of web servers (see, e.g., [9]). These networks also arise in a range of other applications (see, e.g., [22, 7, 21, 27, 43]). Work that focuses on time-dependent phase-type distributions can be found in [36, 37] and references therein.

The focal point of the present paper is the rigorous study of certain aspects of the stochastic optimal control of time-inhomogeneous queueing networks. In most cases, an exact analytic solution is not available. Instead, we use an asymptotic analysis to gain insight into the design of good controls. Specifically, we embed the actual system into a sequence of systems with rates tending to infinity and look for a sequence of controls that are *asymptotically optimal* (in the sense to be described precisely in Definition 3.1).

In many cases, the identification of a class of asymptotically optimal sequences of controls is facilitated by first solving related, but simpler, first-order (or fluid) and/or second-order control problems. The first-order problems arise from functional strong law of large numbers (abbreviated FSLN) limits of the original systems (see,

*Received by the editors June 15, 2009; accepted for publication (in revised form) November 19, 2010; published electronically March 31, 2011. This work was supported in part by the National Science Foundation under grants CMMI-0728064 and CMMI-0928154. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

<http://www.siam.org/journals/sicon/49-2/76202.html>

[†]Department of Mathematics, The University of Texas at Austin, Austin, TX 78712 (mcudina@math.utexas.edu).

[‡]Division of Applied Mathematics, Brown University, Providence, RI 02912 (Kavita_Ramanan@brown.edu).

e.g., Theorem 2.1 of [30]) and lead to deterministic control problems. Second-order problems, additionally, take into account fluctuations around the FLLN limits.

The methodology of using fluid and diffusion control problems to identify asymptotically optimal controls for queueing networks is fairly well developed in the time-homogeneous setting. In the time-homogeneous case, the second-order approximation of a queueing network is usually given by a reflected diffusion, leading to a single reflected diffusion control problem. Historically, asymptotic limit theorems were first established to shed insight into the performance of these networks under various scheduling disciplines. Inspired by these limit theorems, a “formal” limiting control problem was then proposed (say, the Brownian control problem (BCP) for systems in heavy traffic; see, e.g., [49] for references on this subject). Only subsequently were rigorous theorems established in specific settings to link the solution of the limiting control problem to the so-called prelimit control problem (see [4, 5, 6] for examples). Other references in this context include [13, 17, 33, 34, 35] for the use of fluid control problems and [1, 23, 25, 29] for the use of diffusion control problems.

1.1. Time-inhomogeneous networks—Performance analysis. Thus far, the study of time-inhomogeneous networks has mainly concentrated on performance analysis. The seminal paper of Mandelbaum and Massey [30] is the cornerstone of the rigorous approach to the identification of both first-order and second-order approximations for the $M_t/M_t/1$ queue under the *uniform acceleration* regime. The authors of [30] employ the theory of strong approximations (see, e.g., [10, 11, 20]) to develop a Taylor-like expansion of sample paths of queue lengths, establishing an FLLN and a functional central limit theorem (FCLT). Furthermore, explicit forms of the first-order (in the almost sure sense) and second-order (in the distributional sense) approximations of the queue lengths are identified. Chapter 9 of [47] (see also Theorem 3.2 of [31]) relaxes certain technical assumptions imposed in [30] and exhibits more general results. The second-order approximation obtained in [30] can be viewed as a *directional derivative* of the one-sided reflection map (with respect to an appropriate topology on the space of functions). With a view towards establishing analogous approximations for networks with time-inhomogeneous arrival and service rates, properties of directional derivatives of multidimensional reflection maps corresponding to a general class of queueing networks were established in [31]. The article [31] also contains an intuitive introduction into this theory, as well as an overview of related references.

1.2. Time-inhomogeneous networks—Optimal control. In the domain of time-inhomogeneous networks, while heuristics for designing controls were proposed by Newell [41], there is relatively little rigorous work. A noteworthy example of an optimal control problem with a fluid model in the time-inhomogeneous setting is given in [8], where the authors study a particular optimal resource allocation problem for a (stochastic) fluid model with multiple classes, and a controller who dynamically schedules different classes in a system that experiences overload. To the best of the authors’ knowledge, there are no general results in the time-inhomogeneous setting that rigorously show convergence of value functions of the prelimit control problems to the value function of a limiting control problem. As mentioned above, even for the time-homogeneous setting, a general theorem of this nature was obtained only relatively recently [5, 6]. In fact, even a concept akin to the notion of *fluid-scale asymptotic optimality* described in [35] for time-homogeneous networks appears not to have been formulated in the time-inhomogeneous setting.

One of the main aims of this paper is to take a step towards building a useful

framework for the asymptotic optimal control of time-inhomogeneous networks. The specific types of controls we focus on are arrival and/or service rates in time-varying single-class queueing networks. Our goal is to develop a general methodology for the identification of high-performance controls for a time-varying network, based on an optimal control analysis of a related (fluid) approximation. While this philosophy is similar to that used for time-homogeneous networks, the analysis is significantly more involved. To begin with, the nature of the asymptotic approximation has to be modified so as to capture the time-varying behavior. In particular, the so-called uniform acceleration technique is used to embed the particular queueing system into a sequence of systems, which, once properly scaled, converges to a deterministic fluid limit system in the strong sense. We refer to the systems in this sequence with uniformly accelerated rates as the *prelimit systems*. With the view that optimal control problems for the fluid limit are typically more tractable than for the prelimit, we wish to answer the following question:

Can we identify a broad class of performance measures for which the solution of the fluid control problem suffices to construct asymptotically optimal sequences of controls?

The phrase “suffices to construct” above can be interpreted in many ways. For instance, one may resolve to use *exactly* a policy that is optimal for the fluid control problem when controlling the prelimit systems, one may try to express optimal policies for the fluid control problem in terms of a *state dependent (feedback) rule* and then use this rule to control the prelimit systems, or one may opt for a heuristic way to “tweak” optimal controls for the fluid problem to perform well in the prelimit. We choose to focus on (a generalization of) the first of the above-mentioned options. More precisely, we refer to any sequence of policies whose performances in the fluid system converge to the optimum value for the fluid system as a fluid-infimizing sequence, and we seek to identify broad classes of performance measures for which any fluid-infimizing sequence of disciplines is also asymptotically optimal. The main theoretical results of this paper are Theorems 5.4 and 5.5, which provide two alternative sufficient conditions for performance measures to have this property, with the latter condition assuming the existence of an optimal control for the fluid problem.

Although such a connection between the fluid and prelimit control problems may appear intuitive, in section 8.2 we describe several reasonable situations when this fails to hold. This underscores the need for a rigorous analysis to determine when this intuition is indeed valid. We also note that the task of identifying an optimal policy for the fluid control problem is not always straightforward. One approach, exploiting the results of [31] on directional derivatives of the oblique reflection map (ORM), is illustrated in section 7.2.3. This calculus of variations-based technique may be of independent interest.

1.3. Outline of the paper. The paper is organized as follows. The general stochastic optimal control problem of interest is presented in section 2, and the notion of asymptotic optimality is formulated in section 3. The related fluid control problem is described in section 4. The identification question posed above is formalized via the notion of *fluid-optimizability* of performance measures in section 5, where our main results are also stated and proved. Section 6 is dedicated to a description of relevant fluid-optimizable performance measures involving aggregate Lipschitz holding costs. Two examples of network control problems with fluid-optimizable performance measures are presented in section 7. For each example, the fluid control problem

is solved, the solution is used to construct an asymptotically optimal sequence of controls, and simulations that capture the effect of these controls on the original network problem are also presented. Concluding remarks and, in particular, examples where the correspondence between the fluid and original control problems fails to hold, are given in section 8. All auxiliary technical results are gathered in the appendices.

1.4. Notation and technical paraphernalia. The following (standard) notation will be used throughout the paper:

- $T \in (0, \infty)$ denotes the fixed time horizon;
- $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$;
- $\mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ denotes the set of all (\mathbb{P} -a.s.-equivalence classes of) nonnegative random variables on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$;
- $\text{meas}(\cdot)$ denotes Lebesgue measure on \mathbb{R} ;
- $\mathbb{L}^1[0, T]$ denotes the set of all integrable functions defined on $[0, T]$ (with respect to Lebesgue measure);
- $\mathbb{L}_+^1[0, T]$ denotes the set of all nonnegative functions in $\mathbb{L}^1[0, T]$;
- $\mathcal{C}[0, T]$ is the set of all continuous functions $x : [0, T] \rightarrow \mathbb{R}$;
- $\mathcal{I} : (\mathbb{L}^1[0, T])^k \rightarrow (\mathcal{C}[0, T])^k$ is the (integral) mapping defined by

$$\mathcal{I}_t(f) = \left(\int_0^t f^1(s) ds, \dots, \int_0^t f^k(s) ds \right), \quad t \in [0, T],$$

- with $f = (f^1, \dots, f^k) \in (\mathbb{L}^1[0, T])^k$;
- $\mathcal{D}[0, T]$ denotes the set of all real-valued right-continuous functions on $[0, T]$ with finite left limits at all points in $(0, T]$;
- $\mathcal{D}_\uparrow[0, T]$ denotes the subset of $\mathcal{D}[0, T]$ containing all nondecreasing functions;
- $\mathcal{D}_{\uparrow, f}[0, T]$ stands for the subset of $\mathcal{D}_\uparrow[0, T]$ containing functions with at most finitely many jumps;
- $\|\cdot\|_T$, defined by $\|x\|_T = \sup_{t \in [0, T]} |x(t)|$ for $x \in \mathcal{D}$, is the uniform convergence norm on the space $\mathcal{D}[0, T]$;
- $\mathcal{B}(Y)$ denotes the Borel σ -algebra on the topological space Y .

For the sake of completeness, we provide the following definitions to be used in what follows.

DEFINITION 1.1. *Let $R \in \mathbb{R}^{\kappa \times \kappa}$ have positive diagonal elements, and let x be in $(\mathcal{D}[0, T])^\kappa$. We say that a pair $(z, l) \in (\mathcal{D}[0, T])^\kappa \times (\mathcal{D}_\uparrow[0, T])^\kappa$ solves the oblique reflection problem (ORP) associated with the constraint matrix R for the function x if $x(0) = z(0)$ and if for every $t \in [0, T]$,*

- (i) $z(t) \geq 0$;
- (ii) $z(t) = x(t) + Rl(t)$;
- (iii) $\int_0^t \mathbf{1}_{[z^i(s) > 0]} dl^i(s) = 0$ for $i = 1, \dots, \kappa$.

Given a matrix R , if for every $x \in (\mathcal{D}[0, T])^\kappa$ there exists a unique pair (z, l) as above, we define the oblique reflection map $\Gamma : (\mathcal{D}[0, T])^\kappa \rightarrow (\mathcal{D}[0, T])^\kappa$ as $\Gamma(x) = z$ for every $x \in (\mathcal{D}[0, T])^\kappa$.

Motivated by the study of single-class open queueing networks, for a particular class of matrices R , existence and uniqueness of the ORM on the space of continuous functions was proved in the seminal paper [24]. The results in [24] can be directly extended to càdlàg functions to support the above definition (see, for example, Theorem 2.1 in [15]). Generalizations of the ORM that are useful for more general networks are considered in [2, 14, 15, 16, 42] and references therein. Explicit formulas for the solution map associated with certain one-dimensional problems can be found in [28] and [46].

2. Optimal control of time-inhomogeneous, single-class queueing networks. The main goal of the present paper is to elucidate the relationship between fluid optimality and asymptotic optimality (both to be defined precisely in what follows) in the case of single-class, open, time-varying queueing networks with a fixed finite number κ of stations (nodes) and fixed routing dynamics, operated under the FIFO service discipline. The primitive data and dynamical equations governing such a network are introduced in sections 2.1 and 2.2. The class of performance measures under consideration is described in section 2.3.

2.1. Primitive data. Assuming that each station is initially empty and has infinite waiting room, the dynamics of any such network are determined by a pair of processes (\mathbf{E}, \mathbf{S}) defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where

- $\mathbf{E} = (E^i, i = 1, 2, \dots, \kappa) \in (\mathcal{D}_{\uparrow, f}[0, T])^\kappa$ stands for the vector of (cumulative) exogenous arrivals to each of the κ stations;
- $\mathbf{S} = (S^{i,j}, i = 1, 2, \dots, \kappa, j = 1, 2, \dots, \kappa + 1) \in (\mathcal{D}_{\uparrow, f}[0, T])^{\kappa^2 + \kappa}$ denotes the $\kappa \times (\kappa + 1)$ matrix of (cumulative) *potential* service completions in the κ stations; i.e., for all pairs of indices $(i, j) \in \{1, 2, \dots, \kappa\}^2$ the entry $S^{i,j}$ in the matrix stands for the process of (cumulative) potential services at the i th station that would be routed to the j th station and, for $i = 1, \dots, \kappa$, $S^{i, \kappa+1}$ represents the (cumulative) number of jobs that would complete service at station i and leave the network if the i th station were always busy.

In this paper, we focus on the case when \mathbf{E} and \mathbf{S} are constructed from Poisson point processes (PPPs), with rates determined by the functions $\boldsymbol{\lambda} = (\lambda^1, \dots, \lambda^\kappa) \in (\mathbb{L}_+^1[0, T])^\kappa$ and $\boldsymbol{\mu} = (\mu^1, \dots, \mu^\kappa) \in (\mathbb{L}_+^1[0, T])^\kappa$, and a “routing” matrix $P = (p_{ij}; 1 \leq i, j \leq \kappa)$ in the manner described below. For a thorough and concise treatment of PPPs, the reader should consult [26]. The component functions of $\boldsymbol{\lambda}$ represent the time-varying rates of exogenous arrivals to their respective nodes, while the component functions of $\boldsymbol{\mu}$ correspond to the rates of potential services at each of the κ stations. Transitions from a station i to another station j are not deterministic; they are governed by the probabilities encoded in the matrix $P = (p_{ij}; 1 \leq i, j \leq \kappa)$ as follows: once a job is completed at the i th station, it queues up at the j th station with probability p_{ij} . The job leaves the network altogether with probability $1 - \sum_{j=1}^\kappa p_{ij}$. We assume that matrix $P \in \mathbb{R}^{\kappa \times \kappa}$ has spectral radius strictly less than 1.

Remark 2.1. The above condition on P implies that the constraint matrix $R = I - P^T$ associated with the routing matrix P , in the sense of Remark 1.1 of [31], satisfies the [H-R] condition of Definition 1.2 of [31]. This assumption ensures that the ORP associated with R admits a unique solution and that the associated reflection mapping, which we denote by Γ^P , is Lipschitz continuous. For more details, the reader is directed to Theorem 3.1 of [31].

Suppose the primitive data $(\boldsymbol{\lambda}, \boldsymbol{\mu}, P)$ are given, and let $\zeta = (\zeta^1, \dots, \zeta^\kappa)$ and $\xi = (\xi^1, \dots, \xi^\kappa)$ be independent vectors of mutually independent PPPs on the domains $\mathcal{S} := [0, T] \times [0, \infty)$ and $\mathcal{S}' := [0, T] \times [0, \infty) \times [0, 1]$, respectively, with mean intensity measures $dt \times dx$ and $dt \times dx \times dy$. For each $k \in \{1, 2, \dots, \kappa\}$, the process of exogenous arrivals to the k th station is given by

$$(2.1) \quad E_t^k = E_t^k(\boldsymbol{\lambda}) \doteq \zeta^k \{ (s, x) : s \leq t, x \leq \lambda^k(s) \} \quad \text{for every } t \in [0, T].$$

Analogously, we model the potential service process at the k th station representing the jobs that would transition on completion into the j th station as

$$(2.2) \quad S_t^{k,j} = S_t^{k,j}(\boldsymbol{\mu}) \doteq \xi^k \left\{ (s, x, y) : s \leq t, x \leq \mu^k(s), \sum_{i=1}^{j-1} p_{ki} < y \leq \sum_{i=1}^j p_{ki} \right\},$$

and the jobs that would leave the network as

$$(2.3) \quad S_t^{k,\kappa+1} = S_t^{k,\kappa+1}(\boldsymbol{\mu}) \doteq \xi^k \left\{ (s, x, y) : s \leq t, x \leq \mu^k(s), \sum_{i=1}^{\kappa} p_{ki} < y \leq 1 \right\}$$

for $t \in [0, T]$.

Remark 2.2.

- (i) We assume that the routing matrix P is fixed throughout, and we do not emphasize the dependence of the process \mathbf{S} on P in the notation.
- (ii) Note that the above definitions can be naturally extended to the case of random rates $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ that are defined on the same probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and take values in $(\mathbb{L}_+^1[0, T])^{2\kappa}$. We will need this extension in what follows.

2.2. Dynamic equations. We now show how the evolution of the network model can be uniquely determined from the primitive data $(\boldsymbol{\lambda}, \boldsymbol{\mu}, P)$ and associated processes (\mathbf{E}, \mathbf{S}) . Consider the following system of equations:

$$(2.4) \quad \begin{aligned} \tilde{S}_t^{k,j} &= \xi^k \left\{ (s, x, y) : s \in B_t^k, x \leq \mu^k(s), \sum_{i=1}^{j-1} p_{ki} < y \leq \sum_{i=1}^j p_{ki} \right\}, \\ B_t^k &= \{s \leq t : Z_s^k > 0\}, \\ Z_t^k &= E_t^k + \sum_{j=1}^{\kappa} \tilde{S}_t^{j,k} - \sum_{j=1}^{\kappa+1} \tilde{S}_t^{k,j}, \quad k = 1, \dots, \kappa, t \in [0, T], \end{aligned}$$

where

- $B_t = (B_t^1, \dots, B_t^\kappa)$ is a vector of stochastic processes on $[0, T]$ with values in $\mathcal{B}([0, T])$; for every k and t , the set B_t^k stands for the period up to time t during which the k th queue in the system was not empty;
- $\tilde{\mathbf{S}}_t = (\tilde{S}_t^{k,j}, 1 \leq k \leq \kappa, 1 \leq j \leq \kappa + 1) \in (\mathcal{D}_{\uparrow, f}[0, T])^{\kappa+\kappa^2}$ denotes the matrix of random processes of *actual* service completions in the κ stations indexed by k , depending on whether they depart to a station $j = 1, \dots, \kappa$ or leave the network (for $j = \kappa + 1$);
- $\mathbf{Z}_t = (Z_t^k, k = 1, 2, \dots, \kappa) \in (\mathcal{D}[0, T])^\kappa$ stands for the vector of queue-length processes in the κ stations.

It can be shown that system (2.4) uniquely describes the dynamics of an open network using the principle of mathematical induction on the number of stopping times representing the times of arrivals or potential departures from the stations. With probability one, there are at most a finite number of such events during the time interval $[0, T]$ because the stochastic processes modeling the times of these events are PPPs; hence, the principle of mathematical induction is applicable. Recalling that all the PPPs above are assumed to be mutually independent, with probability one there are no two stopping times in the inductive scheme that happen simultaneously. So, the resulting solution to the system (2.4) is unique with probability one. It is worth noting that the above construction departs from the one commonly used for time-homogeneous systems. Here, one keeps track of the entire set of times when a station is empty and loss of service is possible, and not only of the length of that time.

Moreover, in the case of a feedforward network (i.e., for P being a strictly upper-triangular 0-1 matrix), the progression of completed jobs through the system becomes deterministic. So, \mathbf{Z} admits an alternative representation in terms of the so-called

netput process $\mathbf{X} = (X^{(1)}, \dots, X^{(\kappa)})$, which is defined by

$$(2.5) \quad X_t^i \doteq E_t^i - \sum_{j=1}^{\kappa+1} S_t^{i,j} + \sum_{j=1}^{\kappa} S_t^{j,i}, \quad t \in [0, T], i = 1, \dots, \kappa.$$

The routing matrix P associated with a feedforward network has spectral radius zero. Hence, the associated constraint matrix $R = I - P$ satisfies the [H-R] condition (see Definition 1.2 of [31]), and Theorem 3.1 of [31] can be used to show that \mathbf{Z} satisfies

$$(2.6) \quad \mathbf{Z} = \Gamma^P(\mathbf{X}),$$

where Γ^P denotes the multidimensional ORM associated with P , as stated in Remark 2.1.

2.3. The optimal control problem. The performance of a given network can be viewed as a function $J : (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2} \rightarrow \mathbb{R}$ that maps (\mathbf{E}, \mathbf{S}) to the real-valued performance measure of interest. The formal definition, which imposes additional technical conditions, is as follows.

DEFINITION 2.3. Any mapping $J : (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2} \rightarrow \mathbb{R}$ that is bounded from below and Borel measurable, with $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$ endowed with the Borel σ -algebra in the uniform topology, is called a performance measure.

When (\mathbf{E}, \mathbf{S}) are constructed from the primitive data $(\boldsymbol{\lambda}, \boldsymbol{\mu}, P)$ as described in section 2.1, for fixed P , the rates $(\boldsymbol{\lambda}, \boldsymbol{\mu})$ represent the only ingredient of the modeling equations that can (potentially) be varied by the controller. It is reasonable to assume that the controller can observe the system but cannot predict its future behavior. Technically, admissible controls must be *nonanticipating*, i.e., predictable with respect to the filtration $\{\mathcal{H}_t\}$ defined by

$$(2.7) \quad \mathcal{H}_t \doteq \sigma(\zeta(A) : A \in \mathcal{B}([0, t] \times [0, \infty))^\kappa) \vee \sigma(\xi(B) : B \in \mathcal{B}([0, t] \times [0, \infty) \times [0, 1])^\kappa).$$

In addition, we allow for the incorporation of certain exogeneous constraints that may have to be imposed on the set of rates that the controller can choose at any given time. Let \mathbb{A} stand for the subset of $(\mathbb{L}_+^1[0, T])^{2\kappa}$ containing the rates that respect these constraints, and let \mathcal{A} denote the set of all $\{\mathcal{H}_t\}$ -predictable random processes whose trajectories take values in \mathbb{A} . We refer to the set \mathbb{A} as the set of *deterministic admissible controls* and to the set \mathcal{A} as the set of *admissible controls*.

Remark 2.4. The above notion of admissibility implies that the controller has full information of the past and present of a run of the system. This means that the constraints imposed on the admissible control policies are, by construction, quantitative. In this paper, we do not consider optimal control problems that involve constraints based on the *information* available (e.g., cases of delayed information of the state of the system). However, we do address the extreme case of lack of information on the evolution of the system when we consider deterministic (i.e., not state dependent) controls.

For any $(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{A}$, we can define $\mathbf{E}(\boldsymbol{\lambda}) = (E^1(\boldsymbol{\lambda}), \dots, E^\kappa(\boldsymbol{\lambda}))$ and $\mathbf{S}(\boldsymbol{\mu}) = (S^1(\boldsymbol{\mu}), \dots, S^\kappa(\boldsymbol{\mu}))$ via (2.1), (2.2), and (2.3), though now $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ are stochastic (as opposed to deterministic) processes (see Remark 2.2(ii)). It is natural to consider the following control problems: given a performance measure J , identify

$$(2.8) \quad \inf_{(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{A}} J(\mathbf{E}(\boldsymbol{\lambda}), \mathbf{S}(\boldsymbol{\mu})),$$

where the minimum is in the almost sure sense, or identify

$$(2.9) \quad \inf_{(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{A}} \mathbb{E}[J(\mathbf{E}(\boldsymbol{\lambda}), \mathbf{S}(\boldsymbol{\mu}))],$$

assuming the expectation above is well defined, and the associated (sequences of) infimizing controls. A concrete example of such an optimal control problem is provided in section 6. Henceforth, we shall consider the performance measure J as being arbitrary but fixed.

3. Definition of asymptotically optimal controls. Unfortunately, in most situations of interest, the control problems introduced in (2.8) and (2.9) are not explicitly solvable. Instead, in this section, we consider a sequence of “uniformly accelerated” systems and study the related problem of identifying an asymptotically optimal sequence of controls (in the sense of Definition 3.1 below). As will be shown in section 5, for a large class of performance measures that satisfy a certain continuity condition, this problem reduces to the (typically easier) problem of solving a related deterministic optimal control problem (the fluid control problem introduced in section 4). Moreover, as discussed in section 7, the asymptotically optimal sequence of controls can be used to identify near-optimal controls for systems whose parameters lie in the appropriate asymptotic regime.

Let \mathcal{A} be the set of admissible controls defined in section 2.3. With any given routing matrix P , we associate a sequence of performance measures $\{J_n\}_n$ corresponding to a sequence of networks with routing matrix P and with “uniformly accelerated” rates. More precisely, we define the mapping $J_n : \mathcal{A} \rightarrow \mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ by

$$(3.1) \quad J_n(\boldsymbol{\lambda}, \boldsymbol{\mu}) \doteq J\left(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu})\right) \quad \text{for every } (\boldsymbol{\lambda}, \boldsymbol{\mu}) \text{ in } \mathcal{A},$$

with $\mathbf{E}(n\boldsymbol{\lambda})$ and $\mathbf{S}(n\boldsymbol{\mu})$ defined as in (2.1), (2.2), and (2.3). Given a performance measure J and a resulting sequence $\{J_n\}_n$ of performance measures associated with a sequence of uniformly accelerated systems, as defined in (3.1), we loosely formulate an *asymptotically optimal control problem* as follows:

Identify an admissible sequence whose performance, in the limit, is no worse than the performance of any other admissible sequence.

Here, an admissible sequence of controls refers to an element of the space $\mathcal{A}^{\mathbb{N}}$ of all sequences of admissible controls. The following definition formalizes the meaning of the solution of the asymptotically optimal control problem loosely posed above.

DEFINITION 3.1. *We say that an admissible sequence $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ in $\mathcal{A}^{\mathbb{N}}$ is*

(i) *asymptotically optimal if*

$$\liminf_{n \rightarrow \infty} [J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)] \geq 0, \mathbb{P}\text{-a.s.} \quad \text{for all } \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}};$$

(ii) *average asymptotically optimal if $\mathbb{E}[J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)] < \infty$ for all $n \in \mathbb{N}$ and*

$$\liminf_{n \rightarrow \infty} \mathbb{E}[J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)] \geq 0 \quad \text{for all } \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}.$$

We now state an elementary, but useful, consequence of Definition 3.1.

LEMMA 3.2. *Let $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ be an admissible sequence.*

(i) *Suppose that $\{J_n^*\}_{n \in \mathbb{N}}$ is a sequence of random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$ such that*

$$(3.2) \quad \liminf_{n \rightarrow \infty} [J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J_n^*] \geq 0, \mathbb{P}\text{-a.s.} \quad \text{for all } \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}$$

and

$$(3.3) \quad \lim_{n \rightarrow \infty} [J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*) - J_n^*] = 0, \mathbb{P}\text{-a.s.}$$

Then $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ is (strongly) asymptotically optimal.

(ii) Suppose that $\{\tilde{J}_n^*\}_{n \in \mathbb{N}}$ is a sequence of real numbers such that

$$\liminf_{n \rightarrow \infty} (\mathbb{E}[J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)] - \tilde{J}_n^*) \geq 0 \quad \text{for all } \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}$$

and

$$\lim_{n \rightarrow \infty} (\mathbb{E}[J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)] - \tilde{J}_n^*) = 0.$$

Then $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ is average asymptotically optimal.

4. A simpler optimal control problem. In section 4.1, we describe a fluid version of the network equations considered in section 2.2. In view of Theorem A.1, the fluid network is the FSLLN limit of a uniformly accelerated sequence of queueing networks. In section 4.2, we present the fluid control problem.

4.1. Fluid network equations. Given a routing matrix P and $(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A}$, a continuous or “fluid” analogue of the network equations introduced in section 2.2 is

$$(4.1) \quad \bar{\mathbf{X}}_t \doteq \mathcal{I}_t(\boldsymbol{\lambda}) - (I - Q)\mathcal{I}_t(\boldsymbol{\mu}), \quad \bar{\mathbf{Z}}_t \doteq \Gamma^P(\bar{\mathbf{X}})_t, \quad t \in [0, T],$$

where

- I denotes the $\kappa \times \kappa$ -dimensional identity matrix;
- $Q = P^\tau$ stands for the transpose of the fixed routing matrix P ;
- $\bar{\mathbf{X}} \in (\mathcal{C}[0, T])^\kappa$ is the vector of mean netput processes in the κ stations;
- $\Gamma^P : (\mathcal{D}_\uparrow[0, T])^\kappa \rightarrow (\mathcal{D}_\uparrow[0, T])^\kappa$ is the ORM generated by the ORP associated with the routing matrix P .

Recall from the definition of $\mathcal{I}(\cdot)$ given in section 1.4 that $\mathcal{I}(\boldsymbol{\lambda})$ and $\mathcal{I}(\boldsymbol{\mu})$ lie in $(\mathcal{C}[0, T])^\kappa$ and, by (2.1), (2.2), and (2.3), represent the mean cumulative exogenous arrivals to and mean cumulative potential service completions in the κ stations, respectively.

4.2. Fluid limit performance measure. Given the definition of the sequence of performance measures $\{J_n\}_{n \in \mathbb{N}}$ in (3.1), the appropriate analogue of the performance measure in the fluid system is the mapping $\bar{J} : \mathbb{A} \rightarrow \mathbb{R}$ given by

$$(4.2) \quad \bar{J}(\boldsymbol{\lambda}, \boldsymbol{\mu}) \doteq J(\mathcal{I}(\boldsymbol{\lambda}), \text{diag}(\mathcal{I}(\boldsymbol{\mu}))\hat{P}) \quad \text{for every } (\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A},$$

where \hat{P} is a $\kappa \times (\kappa + 1)$ matrix obtained by appending the column vector $(1 - \sum_{i=1}^{\kappa} p_{ki}, 1 \leq k \leq \kappa)$ to the routing matrix P , and \mathbb{A} is the set of deterministic admissible controls.

The *fluid control problem* can be formulated as follows: find

$$(4.3) \quad \bar{J}^* \doteq \inf_{(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A}} \bar{J}(\boldsymbol{\lambda}, \boldsymbol{\mu}),$$

and identify an associated sequence of infimizing controls.

DEFINITION 4.1. A sequence $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}} \in \mathbb{A}$ is said to be *fluid infimizing* if

$$\lim_{n \rightarrow \infty} \bar{J}(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*) = \bar{J}^* = \inf_{(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A}} \bar{J}(\boldsymbol{\lambda}, \boldsymbol{\mu}).$$

In particular cases where the optimal value in the fluid control problem is attained, the following definition makes sense.

DEFINITION 4.2. A policy $(\lambda^*, \mu^*) \in \mathbb{A}$ is said to be optimal for the fluid control problem if

$$\bar{J}(\lambda^*, \mu^*) = \bar{J}^* = \inf_{(\lambda, \mu) \in \mathbb{A}} \bar{J}(\lambda, \mu).$$

If an optimal policy for the fluid control problem exists, then a sequence of policies that are identically equal to this optimal policy is clearly fluid infimizing.

5. A criterion for identification of asymptotically optimal controls.

5.1. The notion of fluid-optimizability. The fluid control problem is typically significantly easier to analyze than the original control problems described in (2.8) and (2.9). It is, therefore, natural to pose the following question:

Under what conditions on the performance measure J will fluid-infimizing sequences be (average) asymptotically optimal?

Theorems 5.4 and 5.5 below provide sufficient conditions for an affirmative answer to this question, which is rigorously phrased in terms of the following notion.

DEFINITION 5.1. Let $J : (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ be a performance measure, and let $\{(\lambda_n^*, \mu_n^*)\}_{n \in \mathbb{N}}$ in \mathbb{A} be an associated fluid-infimizing sequence in the sense of Definition 4.1. If

$$(5.1) \quad \liminf_{n \rightarrow \infty} [J_n(\lambda_n, \mu_n) - J_n(\lambda_n^*, \mu_n^*)] \geq 0, \mathbb{P}\text{-a.s.} \quad \text{for all } \{(\lambda_n, \mu_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}},$$

we say that the performance measure J is fluid optimizable. The performance measure J is called average fluid optimizable if

$$(5.2) \quad \liminf_{n \rightarrow \infty} \mathbb{E}[J_n(\lambda_n, \mu_n) - J_n(\lambda_n^*, \mu_n^*)] \geq 0 \quad \text{for all } \{(\lambda_n, \mu_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}.$$

Remark 5.2. Let V_n be the optimal value or performance in the n th system:

$$(5.3) \quad V_n \doteq \inf_{(\lambda, \mu) \in \mathbb{A}} J_n(\lambda, \mu).$$

As shown below, it is almost an immediate consequence of the definition that if J is fluid optimizable, then for any fluid-infimizing sequence $\{(\lambda_n^*, \mu_n^*)\}_{n \in \mathbb{N}}$,

$$(5.4) \quad \lim_{n \rightarrow \infty} [V_n - J_n(\lambda_n^*, \mu_n^*)] = 0.$$

Indeed, by the definition of V_n , for any fluid-infimizing sequence $\{(\lambda_n^*, \mu_n^*)\}_{n \in \mathbb{N}}$,

$$\limsup_{n \rightarrow \infty} [V_n - J_n(\lambda_n^*, \mu_n^*)] \leq 0.$$

On the other hand, for each $n \in \mathbb{N}$, let $(\tilde{\lambda}_n, \tilde{\mu}_n) \in \mathbb{A}$ be such that $V_n \geq J_n(\tilde{\lambda}_n, \tilde{\mu}_n) - 1/n$. Then, because J is fluid optimizable, by (5.1),

$$\liminf_{n \rightarrow \infty} [V_n - J_n(\lambda_n^*, \mu_n^*)] \geq \liminf_{n \rightarrow \infty} \left[J_n(\tilde{\lambda}_n, \tilde{\mu}_n) - \frac{1}{n} - J_n(\lambda_n^*, \mu_n^*) \right] \geq 0.$$

5.2. Main results. For the remainder of the paper, we assume that the constraint set satisfies the following mild assumption.

ASSUMPTION 5.3. *The constraint set \mathbb{A} is bounded in the space $(\mathbb{L}_+^1[0, T])^{2\kappa}$; i.e., there exists a constant M such that, for every $\boldsymbol{\mu} = (\mu^1, \dots, \mu^{2\kappa})$, $\mathcal{I}_T(\boldsymbol{\mu}^k) < M$ for $k = 1, \dots, 2\kappa$.*

We now state the first main result of the paper.

THEOREM 5.4. *If the mapping $J : (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2} \rightarrow \mathbb{R}$ is uniformly continuous with respect to the uniform topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$, then J is fluid optimizable and, \mathbb{P} -a.s.,*

$$(5.5) \quad \lim_{n \rightarrow \infty} J_n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*) = \bar{J}^*.$$

Proof. Let $\{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}}$ be any admissible sequence in \mathbb{A} . We first claim that if J is uniformly continuous with respect to the uniform topology, then \mathbb{P} -a.s.,

$$(5.6) \quad \lim_{n \rightarrow \infty} |J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)| = 0.$$

We defer the proof of the claim and, instead, first show that the theorem follows from this claim. By (5.6) and the definition (4.3) of \bar{J}^* , it follows that

$$\liminf_{n \rightarrow \infty} [J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}^*] = \liminf_{n \rightarrow \infty} [\bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}^*] \geq 0.$$

This shows that condition (3.2) of Lemma 3.2 is satisfied by the constant sequence $J_n^* = \bar{J}^*$, $n \in \mathbb{N}$. On the other hand, let $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ be a fluid-infimizing sequence, and note that, by definition, \mathbb{P} -a.s., $\lim_{n \rightarrow \infty} \bar{J}(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*) = \bar{J}^*$. Together with the fact that (5.6) holds with $\{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}}$ set equal to the sequence $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$, this implies that (5.5) is satisfied. This implies that the constant sequence $\{J_n^* = \bar{J}^*\}_{n \in \mathbb{N}}$ satisfies condition (3.3) and, hence, by Lemma 3.2, it then follows that $\{(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*)\}_{n \in \mathbb{N}}$ is (strongly) asymptotically optimal. This shows that J is fluid optimizable.

It remains only to establish the limit (5.6). Fix an arbitrary $\delta > 0$. Due to the uniform continuity of J , there exists a positive constant $\varepsilon(\delta)$ such that for every $x, y \in (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$, $\|x - y\|_T < \varepsilon(\delta) \Rightarrow |J(x) - J(y)| < \delta$. Hence, for every n , recalling the definitions (3.1) and (4.2) of J_n and \bar{J} , respectively, we see that

$$(5.7) \quad \begin{aligned} & \mathbb{P}[|J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)| > \delta] \\ &= \mathbb{P}\left[\left| J\left(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}_n), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}_n)\right) - J(\mathcal{I}(\boldsymbol{\lambda}_n), \text{diag}(\mathcal{I}(\boldsymbol{\mu}_n)) \cdot \hat{P}) \right| > \delta \right] \\ &\leq \mathbb{P}\left[\left\| \frac{1}{n}(\mathbf{E}(n\boldsymbol{\lambda}_n), \mathbf{S}(n\boldsymbol{\mu}_n)) - (\mathcal{I}(\boldsymbol{\lambda}_n), \text{diag}(\mathcal{I}(\boldsymbol{\mu}_n)) \cdot \hat{P}) \right\|_T > \varepsilon(\delta) \right]. \end{aligned}$$

By Lemma A.2 the arrival and service processes above are all identical in distribution to the Poisson processes with the appropriately chosen rates. So, we can use the submartingale inequality and the expression for the fourth moment of a Poisson random variable (see the proof of Lemma A.3 for a similar procedure) to further bound the last expression in (5.7) by

$$\frac{2\kappa}{n^4\varepsilon(\delta)^4}(3n^2K_n^2 + nK_n),$$

where $K_n = \max_{1 \leq i \leq \kappa} (\mathcal{I}_T((\boldsymbol{\lambda}_n)_i) \vee \mathcal{I}_T((\boldsymbol{\mu}_n)_i))$. Now, due to Assumption 5.3, the sequence $\{K_n\}_{n \in \mathbb{N}}$ is uniformly bounded from above by a constant, and so

$$\sum_{n=1}^{\infty} \mathbb{P}[|J_n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)| > \delta] < \infty.$$

Since the choice of δ was arbitrary, an application of the Borel–Cantelli lemma completes the proof of (5.6) and, hence, of the theorem. \square

When an optimal policy for the fluid control problem exists, we establish an alternative sufficient condition for fluid-optimizability, in which the uniformity assumption is removed but a stronger form of continuity is now required—specifically, with respect to the M'_1 product topology, which is a stronger assumption since the M'_1 topology is weaker than the uniform topology. For the definition of the M'_1 topology as well as a discussion of its basic properties, the reader is referred to section 13.6.2 of [48].

THEOREM 5.5. *Suppose that there exists a policy $(\lambda^*, \mu^*) \in \mathbb{A}$ that is optimal for the fluid control problem and that J is continuous with respect to the uniform topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$. Then, \mathbb{P} -a.s.,*

$$(5.8) \quad \lim_{n \rightarrow \infty} J_n(\lambda^*, \mu^*) = \bar{J}(\lambda^*, \mu^*) = \bar{J}^*.$$

If in addition J is continuous with respect to the product M'_1 topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$, then J is fluid optimizable.

Proof. The FSLLN result established in Theorem A.1 shows that \mathbb{P} -a.s.,

$$\left\| \frac{1}{n} \mathbf{E}(n\lambda^*) - \mathcal{I}(\lambda^*) \right\|_T \rightarrow 0, \quad \left\| \frac{1}{n} \mathbf{S}(n\mu^*) - \text{diag}(\mathcal{I}(\mu^*)) \hat{P} \right\|_T \rightarrow 0.$$

Together with the definitions (3.1) and (4.2) of J_n and \bar{J} , respectively, as well as the continuity of J with respect to the uniform topology, this implies that \mathbb{P} -a.s.,

$$\lim_{n \rightarrow \infty} |J_n(\lambda^*, \mu^*) - \bar{J}(\lambda^*, \mu^*)| = \lim_{n \rightarrow \infty} |J(\frac{1}{n} \mathbf{E}(n\lambda^*), \frac{1}{n} \mathbf{S}(n\mu^*)) - J(\mathcal{I}(\lambda^*), \mathcal{I}(\mu^*))| = 0,$$

and (5.8) follows.

Next, we turn to the proof of the fluid-optimizability of J under the stronger continuity assumption with respect to the M'_1 topology. Fix an ω for which (5.8) holds (all random quantities in the remainder of the proof will be evaluated at that ω without explicit mention). Let $\{(\lambda_n, \mu_n)\}$ be an arbitrary admissible sequence. Then, due to (5.8) and the fact that (λ^*, μ^*) is optimal for the fluid control problem, the left-hand side of (5.1) satisfies

$$(5.9) \quad \begin{aligned} \liminf_{n \rightarrow \infty} [J_n(\lambda_n, \mu_n) - J_n(\lambda^*, \mu^*)] &= \liminf_{n \rightarrow \infty} [J_n(\lambda_n, \mu_n) - \bar{J}(\lambda^*, \mu^*)] \\ &\geq \liminf_{n \rightarrow \infty} [J_n(\lambda_n, \mu_n) - \bar{J}(\lambda_n, \mu_n)]. \end{aligned}$$

Without loss of generality, we can assume that the right-hand side of (5.9) is finite. Indeed, if this assumption is not satisfied, then inequality (5.1) immediately holds. Let $\{(\eta_\ell, \nu_\ell)\}_{\ell \in \mathbb{N}}$ denote the subsequence of pairs $\{(\lambda_{n_\ell}, \mu_{n_\ell})\}_{\ell \in \mathbb{N}}$ along which the limit inferior above is attained as the proper limit. Since Assumption 5.3 implies that the sequence $\{(\eta_\ell, \nu_\ell)\}_{\ell \in \mathbb{N}}$ is uniformly bounded in $(\mathbb{L}_+^1[0, T])^{2\kappa}$, by Lemma B.1 there exists a further subsequence $\{(\eta_{\ell_m}, \nu_{\ell_m})\}_{m \in \mathbb{N}}$ and a function $F \in (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$ such that $(\mathcal{I}(\eta_{\ell_m}), \text{diag}(\mathcal{I}(\nu_{\ell_m})) \hat{P}) \rightarrow F$ as $m \rightarrow \infty$, in the product M'_1 topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$. The assumed continuity of J in the product M'_1 topology and the definition of \bar{J} given in (4.2) yields

$$\lim_{m \rightarrow \infty} \bar{J}(\eta_{\ell_m}, \nu_{\ell_m}) = \lim_{m \rightarrow \infty} J(\mathcal{I}(\eta_{\ell_m}), \text{diag}(\mathcal{I}(\nu_{\ell_m})) \hat{P}) = J(F).$$

On the other hand, by Theorem A.1, the components of the random vector $(\frac{1}{n} \mathbf{E}(n\eta), \frac{1}{n} \mathbf{S}(n\nu))$ converge to the identity function e on $[0, T]$ in the uniform topology. So, we

can utilize Lemma B.3, with $\frac{1}{n_{\ell_m}}E^i(n_{\ell_m}e)$ (respectively, $\frac{1}{n_{\ell_m}}S^{i,j}(n_{\ell_m}e)$) playing the role of Y_n in the lemma, and $\eta_{n_{\ell_m}}^i$ (respectively, $\nu_{\ell_m}^i$) corresponding to the ν_n in the lemma, for $i \in \{1, \dots, \kappa\}$ and $j \in \{1, \dots, \kappa + 1\}$, to conclude that

$$\left(\frac{1}{n_{\ell_m}}\mathbf{E}(n_{\ell_m}\eta_{\ell_m}), \frac{1}{n_{\ell_m}}\mathbf{S}(n_{\ell_m}\nu_{\ell_m})\right) \rightarrow F \quad \text{as } m \rightarrow \infty$$

in the product M'_1 topology. Hence, by the assumed continuity of J ,

$$\lim_{m \rightarrow \infty} J_{n_{\ell_m}}(\eta_{\ell_m}, \nu_{\ell_m}) = J(F).$$

We conclude that the right-hand side of (5.9) is zero, which completes the proof of the fluid-optimizability of J . \square

As an immediate consequence of Theorems 5.4 and 5.5, we have the following.

COROLLARY 5.6. *If either J is uniformly continuous with respect to the uniform topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$, or an optimal policy for the fluid control problem exists and J is continuous with respect to the product M'_1 topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$, then*

$$(5.10) \quad \lim_{n \rightarrow \infty} V_n = \bar{J}^*.$$

If, in addition, J is uniformly bounded, then it is also average fluid optimizable.

Proof. Let $\{(\lambda_n^*, \mu_n^*)\}_{n \in \mathbb{N}}$ be a fluid-infimizing sequence. Under the first set of hypotheses of the corollary, Theorems 5.4 and 5.5 show that J is fluid optimizable. Hence, by (5.4) of Remark 5.2, the limits of V_n and $J_n(\lambda_n^*, \mu_n^*)$, as $n \rightarrow \infty$, coincide. On the other hand, by (5.5) of Theorem 5.4 and (5.8) of Theorem 5.5, the limit of $J_n(\lambda_n^*, \mu_n^*)$ is \bar{J}^* , and thus (5.10) follows.

Now suppose that, in addition, J is uniformly bounded (say, by a finite constant L). Then for any admissible sequence $\{(\lambda_n, \mu_n)\}_{n \in \mathbb{N}} \in \mathbb{A}^n$,

$$J_n(\lambda_n, \mu_n) - J_n(\lambda_n^*, \mu_n^*) = J\left(\frac{1}{n}\mathbf{E}(n\lambda_n), \frac{1}{n}\mathbf{S}(n\mu_n)\right) - J\left(\frac{1}{n}\mathbf{E}(n\lambda_n^*), \frac{1}{n}\mathbf{S}(n\mu_n^*)\right) \geq -2L.$$

Fatou’s lemma and the fluid-optimizability of J (i.e., inequality (5.1)) then show that

$$\liminf_{n \rightarrow \infty} \mathbb{E}[J_n(\lambda_n, \mu_n) - J_n(\lambda_n^*, \mu_n^*)] \geq \mathbb{E}\left[\liminf_{n \rightarrow \infty} (J_n(\lambda_n, \mu_n) - J_n(\lambda_n^*, \mu_n^*))\right] \geq 0,$$

which proves that J is average fluid optimizable. \square

Remark 5.7. It may be intuitive to expect that policies that are optimal for the fluid control problem provide near-optimal policies for the original control problem, at least for some performance measures (see, e.g., [38, 39, 40]). However, the need for a rigorous approach such as the one provided in this paper is underscored by the fact that this may fail to hold in several natural situations. For further discussion of this issue, see section 8.2.

6. A fluid-optimizable performance measure. In this section, we provide an example of an optimal control problem with a performance measure that is fluid optimizable. For the purposes of this example, we assume that the open network has κ stations and a 0-1 strictly upper-triangular routing matrix P (i.e., it is feedforward and with deterministic routing). We consider a performance measure involving so-called holding costs (also referred to as congestion costs) at every station in the network, which are given in terms of nondecreasing functions of the queue lengths, as well as a reward for completion of jobs over a finite interval $[0, T]$. The controller’s goal is to balance the holding cost penalty with the profit generated by the completed

jobs. Cost structures similar to ours arise in inventory control and are quite standard (see, e.g., Chapter 7 of [23] or page 60 of [50]). For more recent examples of similar cost functions, the reader is directed to [18] and [19] (and references therein). Due to the fact that we consider a finite time-horizon $[0, T]$, there is no discounting or time-averaging of the holding cost.

6.1. The performance measure. Let $h^k : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $k = 1, \dots, \kappa$, be locally Lipschitz functions representing the holding costs at the κ stations in the open network. The total holding cost accumulated over the time period $[0, T]$ is

$$(6.1) \quad h(\mathbf{E}, \mathbf{S}) = \sum_{k=1}^{\kappa} \int_0^T h^k(Z_t^k) dt \quad \text{for every } (\mathbf{E}, \mathbf{S}) \in (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2},$$

where $\mathbf{Z} = (Z^1, Z^2, \dots, Z^\kappa)$ is the queue-length vector defined in (2.4). In this context (recall (2.6)), the vector \mathbf{Z} admits the representation

$$(6.2) \quad \mathbf{Z} = \Gamma^P(\mathbf{X}), \quad \mathbf{X} = \mathbf{E} - (I - P^\tau)S,$$

where Γ^P is the ORM associated with the routing matrix P (see Definition 1.1), and $S = (S^k, 1 \leq k \leq \kappa)$ with $S^k = \sum_{i=1}^{\kappa+1} S^{k,i}$.

Let the profit generated by the completion of jobs during the time interval $[0, T]$ be given by a Lipschitz continuous function $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. We introduce the performance measure $J : (\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2} \rightarrow \mathbb{R}$ as

$$(6.3) \quad J(\mathbf{E}, \mathbf{S}) = h(\mathbf{E}, \mathbf{S}) - p \left(\sum_{k=1}^{\kappa} E_T^k - \sum_{k=1}^{\kappa} Z_T^k \right).$$

LEMMA 6.1. *The mapping J defined in equality (6.3) is Lipschitz continuous on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$ with respect to the uniform metric. If, in addition, Assumption 5.3 is satisfied, J is a fluid-optimizable performance measure.*

Proof. Due to the Lipschitz continuity of both the mapping p and the ORM (see Theorem 3.1 of [31]), it suffices to verify the uniform continuity of h .

Consider (\mathbf{E}, \mathbf{S}) and $(\tilde{\mathbf{E}}, \tilde{\mathbf{S}})$ in $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$. Then an application of the triangle inequality yields

$$|h(\mathbf{E}, \mathbf{S}) - h(\tilde{\mathbf{E}}, \tilde{\mathbf{S}})| \leq \sum_{k=1}^{\kappa} \int_0^T \left| h^k(Z_t^k) - h^k(\tilde{Z}_t^k) \right| dt,$$

where $\mathbf{Z} = (Z^1, Z^2, \dots, Z^\kappa)$ and $\tilde{\mathbf{Z}} = (\tilde{Z}^1, \tilde{Z}^2, \dots, \tilde{Z}^\kappa)$ represent the queue-length vectors of (2.4) associated with pairs (\mathbf{E}, \mathbf{S}) and $(\tilde{\mathbf{E}}, \tilde{\mathbf{S}})$, respectively.

For every k and t , due to the Lipschitz continuity of h^k , we have

$$|h^k(Z_t^k) - h^k(\tilde{Z}_t^k)| \leq C^k |Z_t^k - \tilde{Z}_t^k| \leq C^k \|Z^k - \tilde{Z}^k\|_T,$$

where C^k stands for the Lipschitz constant of the mapping h^k . By (6.2) and the Lipschitz continuity of Γ^P (see Theorem 14.3.4 of [48]), we have

$$\|Z^k - \tilde{Z}^k\|_T \leq K(\|\mathbf{E} - \tilde{\mathbf{E}}\|_T \vee \|\mathbf{S} - \tilde{\mathbf{S}}\|_T) \quad \text{for every } k.$$

Combining the last three inequalities, we deduce that the mapping h is indeed Lipschitz and, thus, uniformly continuous with respect to the uniform topology on $(\mathcal{D}_{\uparrow, f}[0, T])^{2\kappa+\kappa^2}$.

In the presence of Assumption 5.3, we invoke Theorem 5.4 to conclude that the performance measure J is fluid optimizable. \square

7. Applications of fluid-optimizability criteria. In this section, we illustrate how the concept of fluid-optimizability can be applied to study two settings of the general network optimal control problem described in section 6. In each setting, the optimal policy for the fluid control problem is explicitly determined and then used to design (near-optimal) controls for a given “prelimit” system. Following the results of section 5, in order to analyze each control problem, we will embed it into a sequence of “uniformly accelerated” systems, with the N th term in the sequence (for some chosen fixed integer N) representing the actual system. The sign $\hat{\cdot}$ used in the parameters, functions, and processes below (see section 7.2.1) refers to the fact that these quantities correspond to the *actual* network control problem of interest. For instance, in the notation used in section 5, the actual arrival rate $\hat{\lambda}$ equals λ_N for an appropriately chosen index N .

In the case of heavy-traffic approximations of time-homogeneous systems, the systems can be indexed by the load ρ , which converges to 1, and so the index to be assigned to any “prelimit” system is automatically determined by the load of that system (for an application of this principle, see section 2.3.1 of [48]). In contrast, for time-inhomogeneous systems approximated using uniform acceleration, there is no natural correspondence between the prelimit system and the index in the sequence. As a result, there are many choices of index possible or, equivalently, many ways in which to embed the actual system of interest into the sequence of uniformly accelerated systems. While the optimal choice of index is an important subject for future research, in sections 7.1.4 and 7.2.4 we show, with the help of simulations, the effect of different choices of the index. The simulations also serve to illustrate the fluid-optimizability result of Theorems 5.4 and 5.5.

7.1. A single-station example.

7.1.1. The optimal control problem. Consider a single station with a given service rate $\mu \in \mathbb{L}_+^1[0, T]$. Using the notation set up in section 6.1, suppose that the strictly increasing, Lipschitz continuous holding cost function h^1 is such that $h^1(0) = 0$, and that the profit function p is the identity function. We wish to minimize J by varying the arrival rate.

Viewing this as a model of inventory control in which there is a storage (holding) cost at the service station but revenue is earned for each product that leaves the system to meet demand, it is natural to assume that the cumulative mean arrivals of materials into a production station do not greatly exceed the available cumulative service. So, we define the constraint set as

$$\mathbb{A} = \{(\lambda^1, \mu^1) \in (\mathbb{L}_+^1[0, T])^2 : \mathcal{I}_T(\lambda^1) \leq 2\mathcal{I}_T(\mu), \mu^1 = \mu\}.$$

7.1.2. The related fluid control problem. As described in section 4.2, the fluid performance measure is

$$\bar{J}(\lambda) = \int_0^T h^1(\bar{Z}_t^1(\lambda)) dt - (\mathcal{I}_T(\lambda) - \bar{Z}_T^1(\lambda)) \quad \text{for every } \lambda \in \mathbb{A},$$

where we suppress the given parameter μ from the notation and set $\bar{X}_t^1(\lambda) = \mathcal{I}_t(\lambda - \mu)$ and $\bar{Z}_t^1(\lambda) = \Gamma(\bar{X}^1(\lambda))_t$ for $\lambda \in \mathbb{L}_+^1[0, T]$, with Γ denoting the reflection map associated with the single queue (i.e., the standard one-sided reflection map on $[0, \infty)$). The fluid control problem consists of minimizing \bar{J} across $\lambda \in \mathbb{A}$.

7.1.3. Solution of the fluid control problem. Intuitively, it appears to be sensible to make the arrival rate match the service rate throughout the time interval $[0, T]$. This maximizes the throughput while leaving the queue empty and the holding cost nil at all times. Indeed, we have the following result.

LEMMA 7.1. *The policy $\lambda^* = \mu$ is fluid optimal for the above fluid control problem.*

Proof. The fluid performance measure \bar{J} admits the following lower bound for every $\lambda \in \mathbb{L}_+^1[0, T]$:

$$\bar{J}(\lambda) = \int_0^T h_1(\bar{Z}_t^1(\lambda)) dt - (\mathcal{I}_T(\lambda) - \bar{Z}_T^1(\lambda)) \geq -\mathcal{I}_T(\lambda) + \bar{X}_T^1(\lambda) = -\mathcal{I}_T(\mu).$$

Since $\lambda^* = \mu$, we obtain

$$\bar{Z}_t^1(\lambda^*) = \bar{X}_t^1(\lambda^*) = \mathcal{I}_t(\lambda^* - \mu) = 0 \quad \text{for all } t \in [0, T].$$

Thus, $\bar{J}(\lambda^*) = -\mathcal{I}_T(\lambda^*) = -\mathcal{I}_T(\mu)$. The policy $\lambda^* = \mu$ attains the lower bound and is, hence, optimal for the fluid control problem. \square

7.1.4. Embedding and simulations. To illustrate the performance when the optimal policy obtained above is used in the prelimit, we ran simulations of the prelimit systems when the optimal policy for the fluid control problem is implemented. All the simulations were conducted in C++, and the graphs were produced by R. We ran simulations of the prelimit systems for a time horizon $T = 1$ and for two choices of the given service rate: the constant service rate $\mu \equiv 1$ and the periodic service rate $\mu(t) = 1 + \sin(10t)$ for $t \in [0, T]$. In both cases, the holding cost function was taken to be the identity. We present the histograms of the costs produced by 1000 simulation runs for these coefficients, along with the sample summary statistics. The histograms of the costs in the case of the constant service rate are depicted in Figure 1, and those in the case of the periodic μ are shown in Figure 2.

As is readily seen, the histogram means are approaching the optimal value for the fluid control problem (recall (5.10) of Corollary 5.6), although this approach appears to be slow (see Remark 7.2 below for more discussion on this issue). It is worth noting that this approach is slightly faster in the case of the constant given service rate, indicating that fluctuations in the rate have an adverse effect on the speed of convergence. It would be of interest to investigate in future work the appropriate notion/measure of fluctuations and its effect on the convergence. Also, the shapes of the histograms are somewhat skewed, possibly as a result of the effect of reflection (the optimal policy for the fluid control problem produces heavy traffic, i.e., constant upward pushing). The summary statistics are collected in Tables 7.1 and 7.2 for constant and periodic service rates, respectively.

Remark 7.2. The approach of the simulated values to the theoretical limiting cost is rather slow. So, we included the results of taking a large uniform acceleration coefficient of 10,000 (see Figure 3). We believe that this is due to the effect of the system being continuously in heavy traffic (under the optimal policy for the fluid control problem). In such situations, the time-mesh should be quite fine because when the uniform acceleration coefficient is large, there is a high probability of an arrival and/or potential departure in any given interval in the time-mesh. Due to the discretization of time, the simulation will set the time of that jump in the simulated process to be the next node in the partition of the interval $[0, T]$. Hence, one needs to be careful to choose a fine enough mesh-size (possibly at the cost of the speed of simulation). We chose the length of every subinterval in the partition to be 10^{-6} .




FIG. 1. *Single station. Histograms of costs realized for the service rate $\mu \equiv 1$ and for uniform acceleration coefficients $n = 50, 100, 1000$. The optimal value of the fluid control problem is -1 .*




FIG. 2. *Single station. Histograms of costs realized for the service rate $\mu(t) = 1 + \sin(10t)$ and for uniform acceleration coefficients $n = 50, 100, 1000$. The optimal value of the fluid control problem is $1.1 - \cos(10) \approx -1.184$.*

TABLE 7.1

Single station. Summary statistics for 1000 simulations of the cost in the case of $\mu \equiv 1$ and for the acceleration coefficients listed in the first column.

Accel. coeff.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
50	-1.2460	-0.9084	-0.8331	-0.8267	-0.7444	-0.4740
100	-1.1380	-0.9373	-0.8821	-0.8834	-0.8276	-0.6391
1000	-1.0380	-0.9815	-0.9640	-0.9632	-0.9457	-0.8605

TABLE 7.2

Single station. Summary statistics for 1000 simulations of the cost in the case of $\mu(t) = 1 + \sin(10t)$ and for the acceleration coefficients listed in the first column.

Accel. coeff.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
50	-1.402	-1.093	-1.007	-1.001	-0.913	-0.517
100	-1.327	-1.127	-1.059	-1.060	-0.993	-0.770
1000	-1.234	-1.166	-1.146	-1.145	-1.125	-1.043
10000	-1.214	-1.191	-1.184	-1.183	-1.176	-1.154

7.2. A tandem queue example. In this section, we look at a control problem involving a tandem queue. We expand the analysis conducted in the above example in the sense that we explicitly describe two possible embeddings of the actual system into the sequence of uniformly accelerated systems and compare the simulation results. Recall that the $\hat{\cdot}$ symbol will be used for the processes, functions, and constants associated with the actual system to be controlled.




FIG. 3. *Single station. Histogram of costs in the optimal control problem aimed at balancing the holding cost and the throughput, realized for $\mu(t) = 1 + \sin(10t)$ and for the uniform acceleration coefficient $n = 10,000$.*

7.2.1. The optimal control problem. Consider a tandem queue, with the processes $\hat{\mathbf{E}}$ and $\hat{\mathbf{S}}$ of exogenous arrivals and potential services at the first and second station, respectively, modeled using PPPs and rates $\hat{\lambda}, \hat{\mu} \in (\mathbb{L}_+^1[0, T])^2$ as in (2.1), (2.2), and (2.3) (with obvious modifications in the notation). For simplicity, we assume that there are no exogenous arrivals to the second station, i.e., that $\hat{\lambda}^2 \equiv 0$. The service rate $\hat{\mu}^1$ in the first station serves as the control, while $\hat{\lambda}^1$ and $\hat{\mu}^2$ are taken to be known (one can assume that $\hat{\lambda}^1$ and $\hat{\mu}^2$ can be estimated through statistics of previous runs of the system). The actual performance measure that we wish to minimize is the aggregate holding cost in both stations, defined by

$$\hat{J}(\hat{\mathbf{E}}, \hat{\mathbf{S}}) = \int_0^T (\hat{h}^1(\hat{Z}_t^1) + \hat{h}^2(\hat{Z}_t^2)) dt$$

for $(\hat{\mathbf{E}}, \hat{\mathbf{S}}) \in (\mathcal{D}_{\uparrow, f}[0, T])^2 \times (\mathcal{D}_{\uparrow, f}[0, T])^6$, where

- \hat{Z}^i denotes the queue length of the i th queue in the tandem for $i = 1, 2$, associated with the arrival and service processes $(\hat{\mathbf{E}}, \hat{\mathbf{S}})$;
- $\hat{h}^1 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is given by $\hat{h}^1(x) = \hat{c}^1 x^2$ for every $x \in \mathbb{R}_+$, with $\hat{c}^1 > 0$ constant;
- $\hat{h}^2 : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is given by $\hat{h}^2(x) = \hat{c}^2 x$ for every $x \in \mathbb{R}_+$ and for a certain constant $0 < \hat{c}^2 < \mathcal{I}_T(\hat{\lambda}^1)$.

Note that \hat{J} is a special case of the performance measure given in section 6 with $p = 0$. The following constraint on $\hat{\mu}^1$ ensures that admissible policies do not have more (mean) cumulative service available than there are (mean) cumulative arrivals:

$$(7.1) \quad \mathcal{I}_T(\hat{\mu}^1) \leq \mathcal{I}_T(\hat{\lambda}^1).$$

Remark 7.3. The above setup can be envisioned as an example of inventory control in a manufacturing system with two phases (one for each station in the tandem

queue) and with separate storage facilities (buffers) at each station at which holding costs corresponding to functions \hat{h}^1 and \hat{h}^2 of the queue lengths are incurred. The controller’s goal is to minimize the total holding cost \hat{J} by varying the service in the first station; the arrivals to the first station can be understood to depend on the arrival of either raw materials or partially completed products from the previous production phase, while the service at the second station could be taken to depend on the demand for the (partially) finished product.

Recalling the construction of a uniformly accelerated sequence of systems described in section 3, let us refer to the arrival and service rates in the first system in the sequence as “basic” and denote them by λ and μ (as in section 3). With an integer N serving as the embedding constant fixed, in order for the actual system to correspond to the N th system in the sequence, the basic arrival rate and the basic service rate μ should be given by $\lambda = \frac{1}{N} \hat{\lambda}$ and $\mu = \frac{1}{N} \hat{\mu}$. Moreover, the performance measure J takes the form

$$(7.2) \quad J(\mathbf{E}, \mathbf{S}) = \int_0^T (h^1(Z_t^1) + h^2(Z_t^2)) dt \quad \text{for } (\mathbf{E}, \mathbf{S}) \in (\mathcal{D}_{\uparrow, f}[0, T])^2 \times (\mathcal{D}_{\uparrow, f}[0, T])^6,$$

with h^i , $i = 1, 2$, given by $h^1(x) = N^2 \hat{c}^1 x^2$ and $h^2(x) = N \hat{c}^2 x$ for every $x \in \mathbb{R}_+$, and where, for $i = 1, 2$, Z^i denotes the queue length of the i th queue in the tandem associated with (\mathbf{E}, \mathbf{S}) (as defined via (2.1), (2.2), and (2.3) for $n = 1$ and the basic arrival and service rates above). Indeed, with these definitions, it is easily seen that $\hat{J}(\hat{E}, \hat{S}) = J_N(\lambda, \mu)$, where J_N is the performance measure of the N th system in the sequence, defined in terms of J via (3.1).

In addition, using the notation introduced in section 2.3, we can translate the constraint (7.1) pertaining to the actual system into the following constraint on the basic controls.

ASSUMPTION 7.4. *The constraint set is*

$$\mathbb{A} = \{(\lambda, \mu) \in (\mathbb{L}_+^1[0, T])^2 \times (\mathbb{L}_+^1[0, T])^2 : \lambda^1, \mu^2 \text{ fixed as above, } \lambda^2 \equiv 0, \mathcal{I}_T(\mu^1) \leq \mathcal{I}_T(\lambda^1)\}.$$

7.2.2. The related fluid control problem. Since h is fluid optimizable, it follows from Definitions 3.1 and 4.2 that to identify a strongly asymptotically optimal sequence for a control problem with h defined in (6.1) as performance measure, it suffices to analyze the corresponding fluid control problem. We illustrate this procedure for the control problem introduced in section 7.2.1, using a calculus of variations type technique that may be more generally applicable.

Consider a fluid tandem queue, with a given deterministic exogenous arrival rate to the first station denoted by $\lambda \in \mathbb{L}_+^1[0, T]$ and a given deterministic service rate in the second station denoted by $\mu^2 \in \mathbb{L}_+^1[0, T]$. Assume that there are no exogenous arrivals to the second station. Our fluid control problem consists of minimizing the aggregate holding cost in this system by varying the service rate μ in the first station across \mathbb{A} . In view of (4.2) and (7.2), we define the fluid-limit holding cost as

$$(7.3) \quad \bar{h}(\mu) = \int_0^T [h^1(\bar{Z}_t^1(\mu)) + h^2(\bar{Z}_t^2(\mu))] dt$$

for every μ such that $\mathcal{I}_T(\mu) \leq \mathcal{I}_T(\lambda)$ with $h^i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $i = 1, 2$, given by $h^1(x) = c^1 x^2$ and $h^2(x) = c^2 x$ for every $x \in \mathbb{R}_+$, where we set $c^1 = N^2 \hat{c}^1$ and $c^2 = N \hat{c}^2$ to simplify the notation, whereas $\bar{Z}^i(\mu)$, $i = 1, 2$, denote the queue lengths in the fluid tandem queue (as a function of μ).

7.2.3. Solution of the fluid control problem. In the present section, we identify a policy that is optimal for the fluid control problem. To keep the calculations as simple as possible and make transparent the illustration of our calculus of variations type of approach to the problem, we additionally set $\hat{\mu}^2 \equiv 0$. The explicit form of the directional derivative of the ORM obtained in [31] plays a crucial role in the calculations. Also, note that the fluid optimal control problem is not trivial. Due to the convexity of the cost structure in the first station, there is a trade-off between the marginal costs in the two stations to be considered. Heuristic considerations may lead one to conjecture that the optimal policy involves a threshold for the queue length in the first station at which the marginal holding cost in the first station starts to exceed the marginal holding cost in the second station. More precisely, because the marginal holding cost in the second station is constant at c^2 and the marginal holding cost in the first station is a linear function of the queue length, we can propose the following policy: do not serve at all in the first station until the marginal holding cost in the first station exceeds the level c^2 (i.e., once the first queue exceeds $C := \frac{c^2}{2c^1}$), and then start matching the arrival rate so that the first queue remains constant at C . We leave it to the interested reader to formalize the above argument. The proof we provide uses, instead, a calculus of variations approach involving directional derivatives, because one of our goals is to propose a method for determining an optimal policy for the fluid control problem that may be more generally applicable in situations where a simple heuristic may not be available to conjecture the form of the optimal control. In order not to impede the flow, we relegate the proof of the following lemma to Appendix C.

LEMMA 7.5. *For the function μ^* , defined by*

$$\mu^* \doteq \lambda \mathbf{1}_{[t_c, T]} \quad \text{with } t_c = \inf \{t \in [0, T] : \mathcal{I}_t(\lambda) > C := \frac{c^2}{2c^1}\},$$

the admissible policy $((\lambda, 0), (\mu^, 0))$ is fluid optimal for the fluid control problem of section 7.2.2.*

Finally, we have the following corollary, which shows that the optimal value for the above fluid control problem does not depend on the embedding constant.

COROLLARY 7.6. *The optimal value for the fluid control problem of section 7.2.2 is given by*

$$\bar{h}(\mu^*) = \hat{c}^1 \int_0^{\hat{t}_c} (\mathcal{I}_t(\hat{\lambda}))^2 dt + \hat{c}^2 \int_{\hat{t}_c}^T \left(\mathcal{I}_t(\hat{\lambda}) - \frac{\hat{C}}{2} \right) dt,$$

with $\hat{t}_c = \inf\{t > 0 : \mathcal{I}_t(\hat{\lambda}) > \hat{C} := \frac{\hat{c}^2}{2\hat{c}^1}\}$ and \hat{c}^1, \hat{c}^2 , and $\hat{\lambda}$ as in section 7.2.1.

Proof. Using the form of the optimal policy for the fluid control problem μ^* obtained in Lemma 7.5, we have that for every $t \in [0, T]$,

$$\bar{Z}_t^1(\mu^*) = \mathcal{I}_t(\lambda) \wedge C \quad \text{and} \quad \bar{Z}_t^2(\mu^*) = (\mathcal{I}_t(\lambda) - C) \vee 0.$$

Hence,

$$\begin{aligned} \bar{h}(\mu^*) &= \int_0^{t_c} c^1 (\mathcal{I}_t(\lambda))^2 dt + \int_{t_c}^T [c^1 C^2 + c^2 (\mathcal{I}_t(\lambda) - C)] dt \\ (7.4) \quad &= c^1 \int_0^{t_c} (\mathcal{I}_t(\lambda))^2 dt + c^2 \int_{t_c}^T \left(\mathcal{I}_t(\lambda) - \frac{C}{2} \right) dt. \end{aligned}$$




FIG. 4. *The tandem queue. Histograms of costs realized for the embedding constant of 50, $\hat{\lambda}(t) = 100(1 + \sin(10t))$, $t \in [0, 1]$, and for uniform acceleration coefficients $n = 50, 500, 1000$. An approximate optimal value for the fluid control problem is 0.190846 (calculated using Mathematica).*

Recalling that $\lambda = \frac{1}{N}\hat{\lambda}$, $c^1 = N^2\hat{c}^1$, and $c^2 = N\hat{c}^2$, we obtain

$$C = \frac{N\hat{c}^2}{2N^2\hat{c}^1} = \frac{\hat{C}}{N},$$

$$t_c = \inf\{t > 0 : \mathcal{I}_t(\lambda) > C\} = \inf\{t > 0 : \mathcal{I}_t(\hat{\lambda}) > \hat{C}\} = \hat{t}_c.$$

With this in mind, the expression for the optimal cost in the fluid control problem of (7.4) takes the form

$$\begin{aligned} \bar{h}(\mu^*) &= N^2\hat{c}^1 \int_0^{\hat{t}_c} \frac{1}{N^2} (\mathcal{I}_t(\hat{\lambda}))^2 dt + N\hat{c}^2 \int_{\hat{t}_c}^T \frac{1}{N} \left(\mathcal{I}_t(\hat{\lambda}) - \frac{\hat{C}}{2} \right) dt \\ &= \hat{c}^1 \int_0^{\hat{t}_c} (\mathcal{I}_t(\hat{\lambda}))^2 dt + \hat{c}^2 \int_{\hat{t}_c}^T \left(\mathcal{I}_t(\hat{\lambda}) - \frac{\hat{C}}{2} \right) dt. \quad \square \end{aligned}$$

7.2.4. Embedding and simulations. Lemmas 6.1 and 7.5, Assumption 7.4, and Theorem 5.4, when combined, show that the sequence of controls constructed from $\hat{\mu}$ is asymptotically optimal. We use this conclusion to design a good control for the system introduced in section 7.2.1.

We set a time-horizon at $T = 1$ and conducted the simulations for the periodic arrival rate $\hat{\lambda}(t) = 100(1 + \sin(10t))$ for $t \in [0, T]$. As in the previous section, the service rate in the second station is set to zero. The constants in the definition of the holding cost function are set to be $\hat{c}^1 = 1/20,000$ and $\hat{c}^2 = 1/200$. We looked at three uniform acceleration coefficients: $n = 50$, $n = 100$, and $n = 1000$. Examining the effect of choosing the embedding constant $N = 50$, we get the fluid performance measure \bar{h} defined in (7.3) with constants $c^1 = 2500\hat{c}^1 = 0.125$ and $c^2 = 50\hat{c}^2 = 0.25$. Using Lemma 7.5, we obtain an optimal control for the fluid control problem of the form $\hat{\mu} = \lambda \mathbf{1}_{[t_c, T]}$ with $t_c = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) \geq 1\}$ and $\lambda(t) = 2(1 + \sin(10t))$. We present the histograms of the costs based on 1000 simulation runs for these coefficients, along with the sample summary statistics. Figure 4 and Table 7.3 summarize the results of applying the optimal policy $\hat{\mu}$ for the fluid control problem to the prelimit systems. The embedding index of 50 corresponds to the actual system in the sense of section 7.2.1, and the outcome of the simulations of the cost of applying the optimal policy for the fluid control problem to the actual system can be seen in the leftmost graph in Figure 4.

Next, we look at the embedding constant $N = 100$ and repeat the simulations described above for uniform acceleration coefficients $n = 50$, $n = 100$, and $n = 1000$. This time, the arrival rates to the first station were either the constant arrival rate

TABLE 7.3

The tandem queue. Summary statistics for 1000 simulations of the cost in case of the embedding constant of 50, $\hat{\lambda}(t) = 100(1+\sin(10t))$, and for the acceleration coefficients listed in the first column.

Accel. coeff.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
50	0.1221	0.1716	0.1926	0.1950	0.2162	0.3370
500	0.1612	0.1846	0.1912	0.1915	0.1982	0.2299
1000	0.1712	0.1862	0.1906	0.1911	0.1959	0.2158




FIG. 5. The tandem queue. Histograms of costs realized for the embedding constant of 100, $\hat{\lambda} \equiv 100$, and for uniform acceleration coefficients $n = 50, 100, 1000$. The optimal value of the fluid control problem is $7/48 \approx 0.14583$.

TABLE 7.4

The tandem queue. Summary statistics for 1000 simulations of the cost in case of the embedding constant of 100, $\hat{\lambda} \equiv 100$, and for the acceleration coefficients listed in the first column.

Accel. coeff.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
50	0.0699	0.1235	0.1436	0.1489	0.1683	0.3724
100	0.0823	0.1291	0.1455	0.1479	0.1642	0.2457
1000	0.1206	0.1405	0.1458	0.1460	0.1516	0.1721

$\hat{\lambda} \equiv 100$ or the periodic arrival rate $\hat{\lambda}(t) = 100(1 + \sin(10t))$, $t \in [0, T]$. The fluid performance measure \bar{h} is again as in (7.3), but with constants $c^1 = c^2 = 0.5$. By Lemma 7.5, the optimal policy for the fluid control problem has the form $\hat{\mu} = \lambda \mathbf{1}_{[t_c, T]}$ with $t_c = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) \geq 1/2\}$ with $\lambda(t) = 1 + \sin(10t)$. The histograms for $n = 50$, $n = 100$, and $n = 1000$ are shown in Figures 5 and 6. The simulated costs of employing the optimal policy for the fluid control problem in the actual system are given in the middle graphs in Figures 5 and 6. The reader interested in comparing the effects of different embedding constants should compare the leftmost graph in Figure 4 to the middle graph in Figure 6. The summary statistics are provided in Tables 7.4 and 7.5.

Figures 7 and 8 show the graphs of the queue lengths as functions of time for a particular simulation with the uniform acceleration factor $n = 1000$ and for constant and periodic arrival rates, respectively. These two figures illustrate the time at which the optimal service for the fluid limit begins in the first station and starts “matching” the arrivals to the first station.

Remark 7.7. Note that the simulation results in the present section indeed illustrate the claims of Theorem 5.4 and Corollary 5.6. In particular, the simulation values become more concentrated around their averages which, in turn, approach the theoretical optimal value for the fluid control problem, which equals the limit of the prelimit optimal values.




FIG. 6. *The tandem queue. Histograms of costs realized for the embedding constant of 100, $\hat{\lambda}(t) = 100(1 + \sin(10t))$, $t \in [0, 1]$, and for uniform acceleration coefficients $n = 50, 100, 1000$. An approximate optimal value of the fluid control problem is 0.190846 (calculated using Mathematica).*

TABLE 7.5

The tandem queue. Summary statistics for 1000 simulations of the cost in case of the embedding constant of 100, $\hat{\lambda}(t) = 100(1 + \sin(10t))$, and for the acceleration coefficients listed in the first column.

Accel. coeff.	Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
50	0.0941	0.1643	0.1929	0.1987	0.2270	0.4049
100	0.1137	0.1720	0.1924	0.1941	0.2141	0.2947
1000	0.1625	0.1840	0.1908	0.1909	0.1972	0.2261




FIG. 7. *The tandem queue. One trajectory of the queue lengths in the first (increasing in the beginning) and second (the other curve) stations for the embedding constant of 100, the uniform acceleration coefficient $n = 1000$, and $\hat{\lambda} \equiv 100$. The time at which service in the first station begins is 0.5.*

8. Concluding remarks and further research. In this section, we note some features we encountered in this work that are unique to the time-inhomogeneous setup. Some of these issues hint at possible directions of future research. Also, we broadly outline particular problems which were among the topics of [12].

8.1. Important distinctions from the time-homogeneous setup. We stress some unique properties of asymptotically optimal control of queueing networks with time-varying rates. We do this by pointing out certain features of optimal control in the time-homogeneous setting (say, the Brownian control problem (BCP) for




FIG. 8. *The tandem queue. One trajectory of the queue lengths in the first (increasing in the beginning) and second (the other curve) stations for the embedding constant of 100, the uniform acceleration coefficient $n = 1000$, and $\hat{\lambda}(t) = 100(1 + \sin(10t))$. The time at which service in the first station begins is approximately 0.3.*

systems in heavy traffic; see, e.g., [49] for references on this subject) and comparing them to the time-inhomogeneous case. In the time-homogeneous context, the only useful option for the control of a given system is the so-called feedback control, i.e., control which observes the system and is dynamically adapted according to the state in which the system is. Also, to accommodate the information available to the controller, a filtration generated by the stochastic processes driving the model of the system at hand (reflected diffusions in the BCP case) is constructed. Both of these issues are illustrated repeatedly throughout the rich literature of optimal control of time-homogeneous networks.

On the other hand, for the asymptotic analysis in the time-inhomogeneous setting, it is possible to consider *deterministic* controls that are prescribed by the controller in advance of the run of the system and which depend only on the given parameters of the model of the system. In fact, optimal policies for the fluid control problem are deterministic, and it is, indeed, sensible to consider their asymptotic optimality (see section 5). Moreover, to allow for stochastic (state dependent) controls, a novel structure of the accumulation of information available to the controller must be formulated to incorporate the past and present of the system. The theory of PPPs proved to be a convenient modeling tool in this respect (see section 2.1). Both of these points are by-products of our analysis of the main problem.

Having proposed an asymptotically optimal sequence, we would like to implement an element of this sequence of controls in the actual system which inspired the problem in the first place. In the case of BCPs, this connection is more or less straightforward (see, e.g., section 5.5 of [48] for an overview). On the other hand, in the case of time-inhomogeneous queues, it is not immediately clear what the appropriate choice of the index of the actual system when embedded in the prelimit sequence of uniformly accelerated systems should be. The question of the choice of this index is not trivial, and we did not attempt to consider it in the present work. However, recalling that the uniform acceleration method preserves the ratio of arrival and service rates, and encouraged by the simulation results presented in section 7 (see also Corollary 7.6), we are hopeful that there is a rich collection of optimal control problems for which

the choice of the index assigned to the actual system will not strongly influence the performance of the class of asymptotically optimal controls constructed. A more rigorous study of this issue would be worthy of future investigation. In the same vein, it would be interesting to construct a “test” model in which it is possible to solve the prelimit stochastic optimal control problems and compare the performance of policies that are optimal for the fluid control problem to the performance of policies that are optimal for the original control problem.

8.2. Pertinent examples in earlier work. In [12], the following points were illustrated:

- not all reasonable performance measures are fluid optimizable;
- even if a performance measure is not fluid optimizable, there may be a substantial family of policies that are optimal for the fluid control problem which yield asymptotically optimal sequences.

To this end, two examples of stochastic optimal control problems were identified—one involving a single station and one involving a tandem queue.

In the single-station example, both the corresponding fluid control problem and the asymptotically optimal control problem were solved. More precisely, a necessary and sufficient condition for optimality in the fluid control problem, as well as a broad class of asymptotically optimal sequences of policies, were identified (see Theorem 3.2.5 (page 40) and Theorem 3.4.8 (page 49), respectively, in [12]). Using these results, it is easy to show that for a certain set of parameters most, but *not all*, optimal policies for the fluid control problem are asymptotically optimal. In addition, it is also possible to construct an example (not studied in [12]) for which there is a unique optimal policy for the fluid control problem that does not generate an asymptotically optimal sequence. All of the above results are easily generalizable to the single station with a feedback loop.

In the tandem queue setup, it was demonstrated that for a certain set of parameters not only is the performance measure in question not average fluid optimizable, but it is not possible to have an asymptotically optimal sequence that consists of deterministic policies (see section 4.7 (page 91) of [12]). This result indicates that in some situations, a first-order analysis may not be sufficient to design near-optimal policies, but a more detailed analysis will be required. This further emphasizes the need for determining rigorous conditions under which a first-order analysis is sufficient.

Appendix A. The functional strong law of large numbers. In this section, we present and prove a version of the FSLLN. We emphasize that this result, albeit very similar in spirit to Theorem 2.1 of [30], is different. Stochastic processes used to model the exogenous arrival and potential service processes in [30] and in the present paper are merely identically distributed. However, since the processes involved are required to converge almost surely, it is necessary to formulate and justify the FSLLN in the present setting. Recall that our model for the primitive processes in the open network via PPPs was necessary to keep track of the accumulation of information available in the associated optimal control problem by means of the filtration $\{\mathcal{H}_t\}$ of (2.7).

THEOREM A.1. *Let $\{\mu_n\}_{n \in \mathbb{N}}$ be a bounded sequence in $\mathbb{L}_+^1[0, T]$, let $p : [0, T] \rightarrow [0, 1]$ be a deterministic measurable function, and let ξ be a PPP on the domain $\mathcal{S} := [0, T] \times [0, \infty) \times [0, 1]$ with Lebesgue measure as the mean intensity measure. Define the sequence of stochastic processes $\{Y^{(n)}\}$ as*

$$Y_t^{(n)} = \xi\{(s, x, y) : s \leq t, x \leq n\mu_n(s), y > p_s\}, \quad t \in [0, \infty), n \in \mathbb{N}.$$

Then, as $n \rightarrow \infty$,

$$(A.1) \quad \left\| \frac{1}{n} Y^{(n)} - \mathcal{I}((1-p)\mu_n) \right\|_T \rightarrow 0, \mathbb{P}\text{-a.s.}$$

To prove this theorem, we start with an equality in distribution. Its proof is straightforward but technical and lengthy. However, since we could not find a reference for the result, we include it here for completeness.

LEMMA A.2. *Suppose that N is a unit Poisson process, and let $\mu \in \mathbb{L}_+^1[0, T]$ and $p : [0, T] \rightarrow [0, 1]$ be deterministic measurable functions. Furthermore, let ξ be a PPP on the domain $\mathcal{S} := [0, T] \times [0, \infty) \times [0, 1]$ with Lebesgue measure as the intensity measure. Define the stochastic process Y as*

$$Y_t = \xi\{(s, x, y) : s \leq t, x \leq \mu_s, y > p_s\}.$$

Then we have the following distributional equality:

$$N(\mathcal{I}((1-p)\mu)) \stackrel{(d)}{=} Y.$$

Proof. Let ζ denote the PPP on $[0, T]$ associated with the Poisson process $N(\mathcal{I}((1-p)\mu))$. On the other hand, consider the point process χ on \mathcal{S} obtained as a ν -randomization of the PPP ξ for the probability kernel ν from \mathcal{S} to $\mathcal{T} := \{0, 1\}$ given by

$$\begin{aligned} \nu((s, x, y), \{1\}) &= \mathbf{1}_{\{x \leq \mu_s, y > p_s\}}, \\ \nu((s, x, y), \{0\}) &= 1 - \nu((s, x, y), \{1\}). \end{aligned}$$

We introduce the point process χ because the point process $\hat{\chi}$ on $[0, T]$, defined as $\hat{\chi}(C) = \chi(C \times [0, \infty) \times [0, 1] \times \{1\})$ on Borel measurable sets $C \subset [0, T]$, is the PPP associated with the process Y . By the uniqueness theorem for Laplace transforms and Lemma 12.1 in [26], the Laplace transform of a point process uniquely determines its law. Hence, it suffices to prove that $\psi_{\hat{\chi}}(f) = \psi_{\zeta}(f)$ for every nonnegative, measurable f , where $\psi_{\hat{\chi}}$ and ψ_{ζ} are the Laplace transforms of point processes $\hat{\chi}$ and ζ , respectively. By Lemma 12.2 from [26], we have that for every nonnegative, Borel measurable $f : \mathcal{S} \times \{0, 1\} \rightarrow \mathbb{R}_+$

$$(A.2) \quad \psi_{\chi}(f) = \mathbb{E}[\exp(\xi(\log(\hat{\nu}(e^{-f})))],$$

where $\hat{\nu}((s, x, y), \cdot) = \delta_{(s, x, y)} \otimes \nu((s, x, y), \cdot)$ for every $(s, x, y) \in \mathcal{S}$. Let us temporarily fix the function f as above and introduce the function $G : \mathcal{S} \rightarrow \mathbb{R}$ as $G = -\log(\hat{\nu}(e^{-f}))$. Using the interpretation of the kernel $\hat{\nu}$ as an operator on the space of measurable functions, the function G can be rewritten more conveniently as

$$\begin{aligned} G(s, x, y) &= -\log \left(\int_{\mathcal{T}} e^{-f((s, x, y), t)} \hat{\nu}((s, x, y), dt) \right) \\ &= -\log \left(\int_{\mathcal{T}} e^{-f((s, x, y), t)} \delta_{(s, x, y)} \otimes \nu((s, x, y), dt) \right) \end{aligned}$$

for every triplet $(s, x, y) \in \mathcal{S}$. The newly introduced function G allows us to rewrite (A.2) as

$$(A.3) \quad \psi_{\chi}(f) = \mathbb{E}[\exp(\xi(\log(\hat{\nu}(e^{-f})))].$$

Directly from the definition, we conclude that G is Borel measurable. Since f is nonnegative, we must have that $e^{-f} \leq 1$, and since ν is a probability kernel, it is necessary that $\hat{\nu}(e^{-f}) \leq 1$. Therefore, $G \geq 0$, and we can use Lemma 12.2 from [26] again to obtain

$$(A.4) \quad \psi_\chi(f) = \mathbb{E}[\exp(-\xi(G))] = \exp\{-\vartheta(1 - e^{-G(s,x,y)})\},$$

where ϑ is the intensity measure of the process ξ , i.e., $\vartheta = \mathbb{E}[\xi]$. Recalling that ξ is a unit PPP on \mathcal{S} , we conclude that

$$(A.5) \quad \psi_\chi(f) = \exp\left\{-\int_{[0,1]} \int_{\mathbb{R}_+} \int_{[0,T]} (1 - e^{-G(s,x,y)}) ds dx dy\right\}.$$

From the definition of G in terms of f , the expression in (A.5) equals

$$\begin{aligned} \psi_\chi(f) &= \exp\left\{-\int_{\mathcal{S}} (1 - e^{\log(\hat{\nu}(e^{-f((s,x,y),\cdot))}))} ds dx dy\right\} \\ &= \exp\left\{-\int_{\mathcal{T}} \left(1 - \int_{\mathcal{T}} e^{-f((s,x,y),t)} \delta_s \otimes \nu((s,x,y), dt)\right) ds dx dy\right\} \\ &= \exp\left\{-\int_{\mathcal{S}} \left(1 - \int_{\mathcal{T}} e^{-f((s,x,y),t)} \nu((s,x,y), dt)\right) ds dx dy\right\} \\ &= \exp\left\{-\int_{\mathcal{S}} \left(1 - e^{-f((s,x,y),1)} \mathbf{1}_{\{x \leq \mu_s, y > p_s\}} - e^{-f((s,x,y),0)} \mathbf{1}_{\{x > \mu_s, y \leq p_s\}}\right) ds dx dy\right\}. \end{aligned}$$

In particular, for all f such that $f(\cdot, 0) = 0$, we have

$$(A.6) \quad \begin{aligned} \psi_\chi(f) &= \exp\left\{-\int_{\mathcal{S}} \left(1 - e^{-f((s,x,y),1)} \mathbf{1}_{\{x \leq \mu_s, y > p_s\}} - (1 - \mathbf{1}_{\{x \leq \mu_s, y > p_s\}})\right) ds dx dy\right\} \\ &= \exp\left\{-\int_{\mathcal{S}} \mathbf{1}_{\{x \leq \mu_s, y > p_s\}} (1 - e^{-f((s,x,y),1)}) ds dx dy\right\}. \end{aligned}$$

Let us define the operator F on real functions on \mathcal{S} to real functions on $\mathcal{S} \times T$ as $F(g)((s,x,y),t) = g(s,x,y) \mathbf{1}_{\{1\}}(t)$. Then we have, using (A.6), that for every measurable $g : \mathcal{S} \rightarrow \mathbb{R}_+$

$$(A.7) \quad \psi_\chi(F(g)) = \exp\left\{-\int_{\mathcal{S}} \mathbf{1}_{\{x \leq \mu_s, y > p_s\}} (1 - e^{-g(s,x,y)}) ds dx dy\right\}.$$

CLAIM A.1. For every Borel measurable $g : \mathcal{S} \rightarrow \mathbb{R}_+$,

$$(A.8) \quad \psi_{\hat{\chi}}(g) = \psi_\chi(F(g)).$$

In order to prove this ancillary claim, we use “measure theoretic induction.”

1° Let g be of the form $g = \mathbf{1}_B$ for a Borel set B in $[0, T]$. Then we have that

$$\psi_{\hat{\chi}}(g) = \mathbb{E}[e^{-\hat{\chi}(g)}] = \mathbb{E}[e^{-\hat{\chi}(B)}].$$

By the definition of $\hat{\chi}$, the above equals

$$\begin{aligned} \psi_{\hat{\chi}}(g) &= \mathbb{E}[e^{-\chi(B \times [0, \infty) \times [0,1] \times \{1\})}] = \mathbb{E}[e^{-\chi(\mathbf{1}_{B \times [0, \infty) \times [0,1] \times \{1\}})}] \\ &= \mathbb{E}[e^{-\chi(\mathbf{1}_{B \times [0, \infty) \times [0,1] \times \tau^{-1} \mathcal{D}[0,T] \times \{1\}})}] = \mathbb{E}[e^{-\chi(F(g))}] = \psi_\chi(F(g)). \end{aligned}$$

2° Let g be a simple function of the form $g = \sum_{m \leq M} c_m \mathbf{1}_{B_m}$, where $\{c_m\}_{m=1}^M$ are positive constants, and the sets $\{B_m\}_{m=1}^M$ are Borel in $[0, T]$ and mutually disjoint. Then the operator F acts on g as

$$\begin{aligned}
 F(g)((s, x, y), t) &= \left(\sum_{m=1}^M c_m \mathbf{1}_{B_m \times [0, \infty) \times [0, 1]}(s, x, y) \right) \mathbf{1}_{\{1\}}(t) \\
 (A.9) \qquad \qquad \qquad &= \sum_{m=1}^M c_m \mathbf{1}_{B_m \times [0, \infty) \times [0, 1]}(s, x, y) \mathbf{1}_{\{1\}}(t) \\
 &= \sum_{m=1}^M c_m F(\mathbf{1}_{B_m \times [0, \infty) \times [0, 1]})(s, x, y).
 \end{aligned}$$

Due to the linearity of the integration with respect to $\hat{\chi}$, we get

$$\psi_{\hat{\chi}}(g) = \mathbb{E} \left[e^{-\sum_{m=1}^M c_m \hat{\chi}(B_m)} \right].$$

By the definition of $\hat{\chi}$, the above equality gives us

$$\psi_{\hat{\chi}}(g) = \mathbb{E} \left[e^{-\sum_{m=1}^M c_m \hat{\chi}(B_m \times [0, \infty) \times [0, 1] \times \{1\})} \right].$$

Finally, using (A.9) and the linearity of χ , we obtain

$$\psi_{\hat{\chi}}(g) = \mathbb{E} \left[e^{-\chi(F(g))} \right] = \psi_{\chi}(F(g)).$$

3° Finally, let $\{g_n\}$ be an increasing sequence of functions satisfying the equality (A.8) and such that $g_n \uparrow g$ pointwise. By the monotone convergence theorem, we have both

$$\psi_{\hat{\chi}}(g) = \lim_{n \rightarrow \infty} \psi_{\hat{\chi}}(g_n) \quad \text{and} \quad \psi_{\chi}(F(g)) = \lim_{n \rightarrow \infty} \psi_{\chi}(F(g_n)).$$

Since functions g_n were chosen so as to satisfy (A.8), the proposed claim (A.8) holds for every appropriate g . We now have that the Laplace transform of the PPP $\hat{\chi}$ acts on nonnegative measurable functions $g : [0, T] \rightarrow \mathbb{R}_+$ in the following way:

$$(A.10) \qquad \psi_{\hat{\chi}}(g) = \exp \left\{ - \int_S \mathbf{1}_{\{x \leq \mu_s, y > p_s\}} (1 - e^{-g(s)}) ds dx dy \right\}.$$

Note that the Laplace transform of the PPP ζ associated with $N(\mathcal{I}((1-p)\mu))$ is given by

$$(A.11) \qquad \psi_{\zeta}(g) = \exp \left\{ - \int_0^T \mu_s (1 - p_s) (1 - e^{-g(s)}) ds \right\}$$

for every Borel measurable $g : [0, T] \rightarrow \mathbb{R}_+$.

CLAIM A.2. For every Borel measurable $g : [0, T] \rightarrow \mathbb{R}_+$,

$$(A.12) \qquad \psi_{\hat{\chi}}(g) = \psi_{\zeta}(g).$$

Using (A.10) and (A.11), the left-hand side in (A.12) becomes

$$\begin{aligned} \psi_{\hat{\chi}}(g) &= \exp \left\{ - \int_S \mathbf{1}_{\{x \leq n\mu_s\}} \mathbf{1}_{\{y > p_s\}} (1 - e^{-g(s)}) ds dx dy \right\} \\ &= \exp \left\{ - \int_0^T \int_0^\infty \mathbf{1}_{\{x \leq n\mu_s\}} \int_0^1 \mathbf{1}_{\{y > p_s\}} (1 - e^{-g(s)}) dx dy ds \right\} \\ &= \exp \left\{ - \int_0^T n\mu_s(1 - p_s)(1 - e^{-g(s)}) ds \right\} = \psi_\zeta(g). \quad \square \end{aligned}$$

We continue with an application of the submartingale inequality.

LEMMA A.3. *For a unit Poisson process N and a sequence $\{\varphi_n\}_{n \in \mathbb{N}}$ bounded in $\mathbb{L}_+^1[0, T]$ we have*

$$\sum_{n=1}^\infty \mathbb{P} \left[\left\| \frac{1}{n} N(n\mathcal{I}(\varphi_n)) - \mathcal{I}(\varphi_n) \right\|_T > \varepsilon \right] < \infty \quad \text{for every } \varepsilon > 0.$$

Proof. For every n , it is readily seen that the process $N(n\mathcal{I}(\varphi_n)) - \mathcal{I}(\varphi_n)$ is a martingale. Thus, we can employ the submartingale inequality to obtain

$$\begin{aligned} \mathbb{P} \left[\left\| \frac{1}{n} N(n\mathcal{I}(\varphi_n)) - \mathcal{I}(\varphi_n) \right\|_T > \varepsilon \right] &= \mathbb{P} \left[\sup_{0 \leq t \leq T} \left(\frac{1}{n} N(n\mathcal{I}_t(\varphi_n)) - \mathcal{I}_t(\varphi_n) \right)^4 > \varepsilon^4 \right] \\ &\leq \frac{\mathbb{E} \left[\left(N(n\mathcal{I}_T(\varphi_n)) - \mathcal{I}_T(\varphi_n) \right)^4 \right]}{n^4 \varepsilon^4} \\ &\leq \frac{3n(\mathcal{I}_T(\varphi_n))^2 + \mathcal{I}_T(\varphi_n)}{n^3 \varepsilon^4}. \end{aligned}$$

The summability of the right-hand side of the above inequality yields the claim of the lemma. \square

The result stated in Theorem A.1 is an easy consequence of Lemmas A.2 and A.3 combined with the Borel–Cantelli lemma.

Appendix B. Auxiliary fluid-optimizability results. For the definitions and the properties of the M_1 and M'_1 topologies, the reader is directed to sections 12.3 and 13.6 of [48], respectively.

LEMMA B.1. *Let the sequence $\{f_n\}_{n \in \mathbb{N}}$ be bounded in $(\mathbb{L}_+^1[0, T])^k$. Then, there exist a function F in $(\mathcal{D}[0, T])^d$ and a subsequence $\{f_{n_k}\}_{k \in \mathbb{N}}$ such that $\mathcal{I}(f_{n_k}) \rightarrow F$ as $k \rightarrow \infty$ in the product M'_1 topology on $(\mathcal{D}[0, T])^d$ and, equivalently, in the weak M_1 topology on $(\mathcal{D}(0, T))^d$.*

Proof. Let $\{q_m\}$ be a sequence containing all rational numbers in the interval $[0, T]$ and the endpoint T . Then, the sequence of d -tuples $\{\mathcal{I}_{q_1}(f_n^1)\}$ (associated with the first term q_1 of the sequence of rational numbers) has a subsequence $\{\mathcal{I}_{q_1}(f_{n_\ell}^1)\}_\ell$ that converges in \mathbb{R} . The sequence $\{\mathcal{I}_{q_1}(f_{n_\ell}^2)\}_\ell$ has a further subsequence that converges in \mathbb{R} . We can continue this construction along the remaining components of the sequence of d -tuples $\{\mathcal{I}_{q_1}(f_n)\}$ to obtain a subsequence that converges in \mathbb{R}^d . A continuation of these constructive steps across the elements of $\{q_m\}$ forms a diagonalization scheme which produces a sequence $\{G_\ell\}$ which is a subsequence of $\{\mathcal{I}(f_n)\}$ and which converges at all the points in the set $(\mathbb{Q} \cap [0, T]) \cup \{T\}$ to a limit in \mathbb{R}^d . We define the function $F : [0, T] \rightarrow \mathbb{R}^d$ by

$$F(r) = \inf_{q \in \mathbb{Q} \cap [r, T]} \lim_{\ell \rightarrow \infty} G_\ell(q).$$

The fact that the component functions of the terms in the sequence $\{G_\ell\}$ are nondecreasing implies that the function F is well defined and that for every $q \in \mathbb{Q} \cap [0, T]$,

$$F(q) = \lim_{\ell \rightarrow \infty} G_\ell(q).$$

Moreover, since F itself has nondecreasing components, all the components of F have both right and left limits at all points in $(0, T)$, the right limit at 0, and the left limit at T . In addition, if necessary, redefining the function F at T as $F(T) = \lim_{t \uparrow T} F(t)$, we can assume that F is left continuous at T . Next, let us extend the component functions of $\{G_\ell\}$ and F to the domain $[0, \infty)$ so that the extensions are linear with the unit slope on $[T, \infty)$. By this construction, we have ensured that the sequence of unbounded, nondecreasing component functions of $\{G_\ell\}$ converges to the corresponding nondecreasing, unbounded component functions of F on a dense subset of $(0, \infty)$. These are precisely the conditions of Theorem 13.6.3 of [48]. So, we conclude that

$$G_\ell \rightarrow F \quad \text{in the product } M'_1 \text{ topology on } (\mathcal{D}[0, \infty))^d.$$

Using the fact that T is a continuity point of F , we can restrict the domain of the functions above and assert that

$$G_\ell \rightarrow F \quad \text{in the product } M'_1 \text{ topology on } (\mathcal{D}[0, T])^d.$$

Since by Theorem 12.5.2 of [48] the weak M_1 topology coincides with the product M_1 topology on $(\mathcal{D}[0, T])^d$, it suffices to again utilize Theorem 13.6.3 of [48] to conclude that the components of $\{G_\ell\}$ converge to the components of F in the M_1 topology on $(\mathcal{D}(0, T))^d$. \square

Remark B.2. We draw attention to the fact that the necessity of the choice of the M'_1 topology in Lemma B.1 stems from the possibility of a jump at 0 of the limiting function F . Unless we either relax the choice of topology from the more conventional M_1 to M'_1 or restrict the domain of the converging subsequences, we can have no hope of obtaining a “relative-compactness-like” result such as the one in Lemma B.1.

We proceed with a simple lemma regarding the convergence in M'_1 of the composition of functions from two particular convergent sequences.

LEMMA B.3. *Let $\{Y_n\}$ and $\{\nu_n\}$ be sequences in $\mathcal{D}_\uparrow[0, T]$ satisfying*

$$\begin{aligned} Y_n &\rightarrow e \quad \text{in the } M'_1 \text{ topology on } \mathcal{D}_\uparrow[0, T] \text{ and} \\ \nu_n &\rightarrow \nu \quad \text{in the } M'_1 \text{ topology on } \mathcal{D}_\uparrow[0, T] \end{aligned}$$

for some function $\nu \in \mathcal{D}_\uparrow[0, T]$ which is left continuous at T and where e denotes the identity function on the interval $[0, T]$. Then, we have

$$Y_n \circ \nu_n \rightarrow \nu \quad \text{in the } M'_1 \text{ topology on } \mathcal{D}_\uparrow[0, T].$$

Proof. It is convenient to reduce the discussion of M'_1 convergence $\mathcal{D}_\uparrow[0, T]$ to the discussion of convergence in the M_1 topology of restrictions of functions in $\mathcal{D}_\uparrow[0, T]$ to $\mathcal{D}_\uparrow[\varepsilon, T]$, $\varepsilon > 0$. To substantiate this statement, recall the manner in which the functions in the proof of Lemma B.1 were extended, and also recall the equivalence relationship of Theorem 13.6.3 in [48] and the fact that all the (linear, increasing extensions of) the functions in the present lemma conform to the conditions outlined

therein. Then, one can invoke the definition of the M_1 topology for functions on noncompact domains from page 414 of [48]. In summary, it suffices to verify that

$$Y_n \circ \nu_n \rightarrow \nu \quad \text{in the } M_1 \text{ topology on } \mathcal{D}_\uparrow[\epsilon, T]$$

for ϵ that are positive continuity points of ν . The last claim is a direct consequence of Theorem 13.2.4 in [48], which completes the proof. \square

Appendix C. The optimal policy for the tandem queue fluid control problem. This section contains the formal proof of Lemma 7.5.

Suppose that a policy that is optimal for the fluid control problem exists and denote it by μ^* . We shall first argue that the following claim holds.

CLAIM C.1. *Without loss of generality, we can assume $\mathcal{I}_t(\mu^*) \leq \mathcal{I}_t(\lambda)$ for all $t \in [0, T]$.*

Proof of Claim C.1. Suppose, to the contrary, that the proposed inequality is violated. The queue lengths in the fluid system must satisfy the equations (4.1). It can be shown that the queue length in the first station, when any $\mu \in \mathbb{L}_+^1[0, T]$ is the service employed there, can be rewritten as

$$(C.1) \quad \bar{Z}_t^1(\mu) = \mathcal{I}_t(\lambda - \mu) + \int_0^t (-\lambda(s) + \mu(s))^+ \mathbf{1}_{[\bar{Z}_s^1(\mu)=0]} ds, \quad t \in [0, T],$$

while the queue length in the second station equals

$$(C.2) \quad \bar{Z}_t^2(\mu) = \mathcal{I}_t(\mu) - \int_0^t (-\lambda(s) + \mu(s))^+ \mathbf{1}_{[\bar{Z}_s^1(\mu)=0]} ds, \quad t \in [0, T].$$

Let us define $\check{\mu} \in \mathbb{A}$ as

$$\check{\mu}(t) = \mu^*(t) - (-\lambda(t) + \mu^*(t))^+ \mathbf{1}_{[\bar{Z}_t^1(\mu^*)=0]}, \quad t \in [0, T].$$

Then, by (C.1), $\bar{Z}^1(\check{\mu}) = \mathcal{I}_t(\lambda) - \mathcal{I}_t(\check{\mu}) = \bar{Z}_t^1(\mu^*) \geq 0$, and by (C.2), $\bar{Z}_t^2(\check{\mu}) = \mathcal{I}_t(\check{\mu}) = \bar{Z}_t^2(\mu^*)$ for every t . Hence, $h(\mu^*) = h(\check{\mu})$, while $\check{\mu}$ satisfies the desired inequality.

Let us return to the proof of the lemma, assuming that μ^* satisfies the inequality in Claim C.1. If μ^* is optimal for the fluid control problem, we refer to every perturbation $\Delta\mu \in \mathbb{L}^1[0, T]$ (which is not necessarily nonnegative or an admissible policy itself) such that $\mu^* + \varepsilon\Delta\mu \in \mathbb{A}$ for all sufficiently small $\varepsilon > 0$ as an admissible perturbation. Then, for every admissible perturbation $\Delta\mu$ and for every constant $\varepsilon > 0$ such that $\mu^* + \varepsilon\Delta\mu \in \mathbb{A}$, we must have

$$(C.3) \quad \bar{h}(\mu^* + \varepsilon\Delta\mu) - \bar{h}(\mu^*) \geq 0.$$

From (C.1) and (C.2) for \bar{Z} , it is clear that for any $\mu \in \mathbb{A}$ that satisfies the condition of Claim C.1,

$$(C.4) \quad \bar{Z}(\mu) = (\bar{Z}^1(\mu), \bar{Z}^2(\mu)) = \Gamma(\bar{X}(\mu)) = \Gamma(\mathcal{I}(\lambda - \mu), \mathcal{I}(\mu - \mu^2)) = \Gamma(\mathcal{I}(\lambda - \mu), \mathcal{I}(\mu)).$$

Therefore, setting $\chi = (\mathcal{I}(-\Delta\mu), \mathcal{I}(\Delta\mu))$, we can write

$$\frac{1}{\varepsilon}(\bar{Z}(\mu^* + \varepsilon\Delta\mu) - \bar{Z}(\mu^*)) = \nabla_\chi^\varepsilon \Gamma(\bar{X}(\mu^*)),$$

where, as in [31], we adopt the notation

$$\nabla_\chi^\varepsilon \Gamma(\psi) \doteq \frac{1}{\varepsilon} [\Gamma(\psi + \varepsilon\chi) - \Gamma(\psi)]$$

for any càdlàg ψ . Using the definition of \bar{h} given in (7.3) and observing that $h^1(x + \Delta x) - h^1(x) = c^1 \Delta x(2x + \Delta x)$ and $h^2(x + \Delta x) - h^2(x) = c^2 \Delta x$, we see that (C.3) holds if and only if

$$(C.5) \quad \frac{1}{\varepsilon} \int_0^T [c^1(\bar{Z}_t^1(\mu^* + \varepsilon \Delta \mu) - \bar{Z}_t^1(\mu^*))(\bar{Z}_t^1(\mu^* + \varepsilon \Delta \mu) + \bar{Z}_t^1(\mu^*)) + c^2(\bar{Z}_t^2(\mu^* + \varepsilon \Delta \mu) - \bar{Z}_t^2(\mu^*))] dt \geq 0.$$

It follows from Theorem 1.1 in [31] that, as $\varepsilon \downarrow 0$, the pointwise limit of $\nabla_\chi^\varepsilon \Gamma(\bar{X}(\mu^*))$ exists and is given explicitly by

$$\nabla_\chi \Gamma(\bar{X}(\mu^*)) = \lim_{\varepsilon \downarrow 0} \nabla_\chi^\varepsilon \Gamma(\bar{X}(\mu^*)) = \chi + (\gamma^1, -\gamma^1 + \gamma^2),$$

where

$$\gamma^1(t) = \sup_{s \in \Phi^1(t)} [\mathcal{I}_s(\Delta \mu)]^+, \quad \gamma^2(t) = \sup_{s \in \Phi^2(t)} [-\mathcal{I}_s(\Delta \mu) + \gamma^1(s)]^+,$$

and $\Phi^i(t) \doteq \{s \leq t : \bar{Z}_s^i(\mu^*) = 0\} = \{s \leq t : \bar{X}_s^i(\mu^*) = 0\}$, $i = 1, 2$, for every $t \in [0, T]$. Here, the latter equality for the sets $\Phi^i(t)$, $i = 1, 2$, $t \in [0, T]$, is implied by the properties of the ORM (see Theorem 1.1 in [31]) and Claim C.1.

Therefore, for every admissible perturbation $\Delta \mu$, taking limits as $\varepsilon \downarrow 0$ in (C.5), we see that

$$(C.6) \quad \int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathcal{I}_t(\Delta \mu) + \gamma^1(t)) + C\gamma^2(t)] dt \geq 0,$$

with $C := \frac{c^2}{2c^1}$. Define the time instances

$$t_0 = \inf\{t \in [0, T] : \mathcal{I}_t(\mu^*) > 0\} \quad \text{and} \quad t_c = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) > C\}.$$

Due to the assumption that $\mu_2 \equiv 0$, and Claim A.1, we immediately conclude that $\Phi^2(t) = [0, t \wedge t_0]$ for every $t \in [0, T]$. We now claim the following.

CLAIM C.2. $t_c \leq t_0$.

Proof of Claim C.2. Consider an admissible perturbation $\Delta \mu$ such that $\Delta \mu(t) \leq 0$ for all $t \in [0, T]$ and $\Delta \mu(t) = 0$ for every $t \in [0, t_0]$. Then, $\gamma^1 \equiv 0$ and $\gamma^2(t) = \sup_{s \in [0, t \wedge t_0]} [-\mathcal{I}_s(\Delta \mu)] = 0$. Therefore, (C.6) reduces to

$$\int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathcal{I}_t(\Delta \mu))] dt = \int_{t_0}^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathcal{I}_t(\Delta \mu))] dt \geq 0.$$

Since $-\mathcal{I}_t(\Delta \mu) \geq 0$ and the above inequality must hold for all such $\Delta \mu$, we conclude that $\bar{Z}_t^1(\mu^*) = \mathcal{I}_t(\lambda - \mu^*) \geq C$ for all $t \geq t_0$, which establishes the claim.

We now show that, in fact, the following holds.

CLAIM C.3. $t_c = t_0$.

Proof of Claim C.3. Let us assume that $t_c < t_0$ and consider an arbitrary admissible perturbation $\Delta \mu \geq 0$. Then, $\gamma^1(t) = \mathcal{I}_{m^1(t)}(\Delta \mu)$, where $m^1(t) := \sup \Phi^1(t)$ for every $t \in [0, T]$. So,

$$\gamma^2(t) = \sup_{s \in \Phi^2(t)} [-\mathcal{I}_s(\Delta \mu) + \gamma^1(s)]^+ = \sup_{s \in \Phi^2(t)} [-\mathcal{I}_s(\Delta \mu) + \mathcal{I}_{m^1(s)}(\Delta \mu)]^+.$$

By definition, $m^1(s) \leq s$ for every s , and so, recalling that $\Delta\mu \geq 0$, we conclude that $\gamma^2 \equiv 0$. Thus, the inequality (C.6) can be rewritten as

$$\int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathcal{I}_t(\Delta\mu) + \mathcal{I}_{m^1(t)}(\Delta\mu))] dt \geq 0$$

for every admissible perturbation $\Delta\mu \geq 0$. Therefore,

$$\int_0^T \int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathbf{1}_{[m^1(t),t]}(u))\Delta\mu(u)] du dt \geq 0$$

for every admissible perturbation $\Delta\mu \geq 0$. Due to Fubini's theorem, the above inequality yields

$$\int_0^T \Delta\mu(u) \left(\int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathbf{1}_{[m^1(t),t]}(u))] dt \right) du \geq 0$$

for every admissible perturbation $\Delta\mu \geq 0$. We define the function $F : [0, T] \rightarrow \mathbb{R}_+$ as

$$F(u) = \int_0^T [(\bar{Z}_t^1(\mu^*) - C)(-\mathbf{1}_{[m^1(t),t]}(u))] dt$$

and deduce that $F(u) \geq 0$, u -a.e. However, for every $u \in (t_c, t_0)$, using Claim C.2, the fact that $-\mathbf{1}_{[0,t_c]}(u) = -\mathbf{1}_{[t_0,T]}(u) = 0$, and $\mathcal{I}_t(\lambda) > C$ for every $t > t_c$, we have

$$F(u) = \int_0^{t_0} [(\bar{Z}_t^1(\mu^*) - C)(-\mathbf{1}_{[0,t]}(u))] dt = \int_{t_c}^{t_0} [(\mathcal{I}_t(\lambda) - C)(-\mathbf{1}_{[0,t]}(u))] dt \leq 0.$$

This leads to a contradiction, and so Claim C.3 follows.

To conclude the proof of the lemma, it suffices to show the next claim.

CLAIM C.4. $\mu^*(t) = \lambda(t)$ for almost every $t \geq t_0$.

Proof of Claim C.4. Let us assume that there exists a pair of time instances $t_1 < t_2$ such that $t_0 < t_1$ and $\bar{Z}_t^1(\mu^*) > C$ for every $t \in (t_1, t_2)$. Note that $\Phi^1(t) \subseteq [0, t_0] \subset (t_1, t_2)^c$ and recall that $\Phi^2(t) = [0, t \wedge t_0]$ for every t . Consider any admissible perturbation $\Delta\mu$ such that $\Delta\mu(t) = 0$ for every $t \in (t_1, t_2)^c$, $\mathcal{I}_{t_1}(\Delta\mu) = \mathcal{I}_{t_2}(\Delta\mu) = 0$, and $\mathcal{I}_t(\Delta\mu) > 0$ for every $t \in (t_1, t_2)$. Then, for such a function $\Delta\mu$, $\gamma^1(t) = \gamma^2(t) = 0$ for all t . Thus, the left-hand side of the inequality (C.6) reads as

$$\int_{t_1}^{t_2} [(\bar{Z}_t^1(\mu^*) - C)(-\mathcal{I}_t(\Delta\mu))] dt.$$

From the choice of $\Delta\mu$ and the definition of t_1 and t_2 , we conclude that the above expression must be strictly negative, which contradicts the inequality (C.6). Thus, $\bar{Z}_t^1(\mu^*) \leq C$ for every $t \in (t_0, T)$.

On the other hand, it has already been shown (see the last line of the proof of Claim C.2) that $\bar{Z}_t^1(\mu^*) \geq C$ for every $t \in (t_0, T)$. Combining the above two inequalities, we conclude that $\bar{Z}_t^1(\mu^*) = C$ for $t \in (t_0, T)$. So, $\mu^*(t) = \lambda(t)$ for almost every $t \in (t_0, T)$. This proves Claim C.4 and, thus, concludes the proof of the lemma. \square

Acknowledgments. The authors owe a debt of gratitude to the anonymous referee whose insights have improved the quality of the paper and its presentation. Both authors would like to thank Steven Shreve for his generous support of the work presented in this paper. The first author also wishes to thank Gordan Žitković for assistance with the simulations presented in this paper and for many invaluable discussions on the results.

REFERENCES

- [1] B. ATA, J. M. HARRISON, AND L. A. SHEPP, *Drift rate control of a Brownian processing system*, Ann. Appl. Probab., 15 (2005), pp. 1145–1160.
- [2] R. ATAR, A. BUDHIRAJA, AND K. RAMANAN, *Deterministic and stochastic differential inclusions with multiple surfaces of discontinuity*, Probab. Theory Related Fields, 142 (2008), pp. 249–283.
- [3] N. BANSAL, T. KIMBREL, AND K. PRUHS, *Dynamic speed scaling to manage energy and temperature*, in Proceedings of the 45th IEEE Symposium on Foundations of Computer Science, 2004, pp. 520–529.
- [4] N. S. L. BELL AND R. J. WILLIAMS, *Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: Asymptotic optimality of a threshold policy*, Ann. Appl. Probab., 11 (2001), pp. 608–649.
- [5] A. BUDHIRAJA AND A. P. GHOSH, *Diffusion approximations for controlled stochastic networks: An asymptotic bound for the value function*, Ann. Appl. Probab., 16 (2006), pp. 1962–2006.
- [6] A. BUDHIRAJA, A. P. GHOSH, AND C. LEE, *An Ergodic Rate Control Problem for Single Class Queueing Networks*, preprint, Iowa State University, Ames, IA, 2008; also available online at aphghosh.public.iastate.edu/ergodic-control-BGL-July302008.pdf
- [7] M. J. CARRILLO, *Extensions of Palm’s theorem: A review*, Management Sci., 37 (1991), pp. 739–744.
- [8] J. CHANG, H. AYHAN, J. G. DAI, AND C. H. XIA, *Dynamic scheduling of a multiclass fluid model with transient overload*, Queueing Syst., 48 (2004), pp. 263–307.
- [9] Y. CHEN, A. DAS, N. GAUTAM, Q. WANG, AND A. SIVASUBRAMANIAM, *Pricing-based strategies for autonomic control of web servers with time-varying request arrivals*, J. Engrg. Appl. Artificial Intell., 17 (2004), pp. 841–854.
- [10] M. CSÖRGÓ AND L. HORVÁTH, *Weighted Approximations in Probability and Statistics*, Wiley Ser. Probab. Math. Statist., John Wiley & Sons, Chichester, UK, 1993.
- [11] M. CSÖRGÓ AND P. RÉVÉSZ, *Strong Approximations in Probability and Statistics*, Probab. Math. Statist., Academic Press [Harcourt Brace Jovanovich Publishers], New York, 1981.
- [12] M. ČUDINA, *Asymptotically Optimal Control for Some Time-Varying Stochastic Networks*, Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA, 2006; also available online at <http://www.math.utexas.edu/users/mcudina/thesis.pdf>.
- [13] P. DUPUIS, *Explicit solution to a robust queueing control problem*, SIAM J. Control Optim., 42 (2003), pp. 1854–1875.
- [14] P. DUPUIS AND K. RAMANAN, *Convex duality and the Skorokhod problem. I*, Probab. Theory Related Fields, 115 (1999), pp. 153–195.
- [15] P. DUPUIS AND K. RAMANAN, *Convex duality and the Skorokhod problem. II*, Probab. Theory Related Fields, 115 (1999), pp. 197–236.
- [16] P. DUPUIS AND K. RAMANAN, *A multiclass feedback queueing network with a regular Skorokhod problem*, Queueing Syst., 36 (2000), pp. 327–349.
- [17] P. DUPUIS AND K. RAMANAN, *An explicit formula for the solution of certain optimal control problems on domains with corners*, Theory Probab. Math. Statist., 63 (2000), pp. 32–48.
- [18] J. M. GEORGE AND J. M. HARRISON, *Dynamic control of a queue with adjustable service rate*, Oper. Res., 49 (2001), pp. 720–731.
- [19] A. GHOSH, S. M. RYAN, L. WANG, AND A. WEERASIGNHE, *Heavy Traffic Analysis of a Simple Closed Loop Supply Chain*, preprint, Iowa State University, Ames, IA, 2009; also available online at smryan.public.iastate.edu/GRWW10.pdf
- [20] P. W. GLYNN, *Strong approximations in queueing theory*, in Asymptotic Methods in Probability and Statistics (Ottawa, ON, 1997), North-Holland, Amsterdam, 1998, pp. 135–150.
- [21] L. V. GREEN AND P. J. KOLESAR, *A note on approximating peak congestion in $M_t/G/\infty$ queues with sinusoidal arrivals*, Management Sci., 44 (1998), pp. 137–144.
- [22] V. GUPTA, M. HARCHOL-BALTER, A. SCHELLER WOLF, AND U. YECHIALI, *Fundamental characteristics of queues with fluctuating load*, in SIGMETRICS ’06/Performance ’06: Proceedings of the Joint International Conference on Measurement and Modeling of Computer Systems, ACM, New York, 2006, pp. 203–215.

- [23] J. M. HARRISON, *Brownian Motion and Stochastic Flow Systems*, Robert E. Krieger Publishing, Malabar, FL, 1990. Reprint of the 1985 original.
- [24] J. M. HARRISON AND M. I. REIMAN, *Reflected Brownian motion on an orthant*, Ann. Probab., 9 (1981), pp. 302–308.
- [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] O. KALLENBERG, *Foundations of Modern Probability*, 2nd ed., Probab. Appl. (N. Y.), Springer-Verlag, New York, 2002.
- [27] C. KNESSL AND Y. P. YANG, *An exact solution for an $m(t)/m(t)/1$ queue with time-dependent arrivals and service*, Queueing Syst. Theory Appl., 40 (2002), pp. 233–245.
- [28] L. KRUK, J. LEHOCZKY, K. RAMANAN, AND S. SHREVE, *An explicit formula for the Skorokhod map on $[0, a]$* , Ann. Probab., 35 (2007), pp. 1740–1768.
- [29] H. J. KUSHNER, *Heavy Traffic Analysis of Controlled Queueing and Communication Networks*, Appl. Math. (N. Y.) 47, Stochastic Modelling and Applied Probability, Springer-Verlag, New York, 2001.
- [30] A. MANDELBAUM AND W. A. MASSEY, *Strong approximations for time-dependent queues*, Math. Oper. Res., 20 (1995), pp. 33–64.
- [31] A. MANDELBAUM AND K. RAMANAN, *Directional derivatives of oblique reflection maps*, Math. Oper. Res., 35 (2010), pp. 527–558.
- [32] W. A. MASSEY, *The analysis of queues with time-varying rates for telecommunications models*, Telecommunication Syst., 21 (2002), pp. 173–204.
- [33] S. P. MEYN, *Sequencing and routing in multiclass queueing networks part I. Feedback regulation*, SIAM J. Control Optim., 40 (2001), pp. 741–776.
- [34] S. P. MEYN, *Sequencing and routing in multiclass queueing networks part II. Workload relaxations*, SIAM J. Control Optim., 42 (2003), pp. 178–217.
- [35] S. MEYN, *Control Techniques for Complex Networks*, Cambridge University Press, Cambridge, UK, 2008.
- [36] B. L. NELSON AND M. R. TAAFFE, *The $Ph_t/Ph_t/\infty$ queueing systems. I. The single node*, INFORMS J. Comput., 16 (2004), pp. 266–274.
- [37] B. L. NELSON AND M. R. TAAFFE, *The $[Ph_t/Ph_t/\infty]^K$ queueing system. II. The multiclass network*, INFORMS J. Comput., 16 (2004), pp. 275–283.
- [38] G. F. NEWELL, *Queues with time-dependent arrival rates. I. The transition through saturation*, J. Appl. Probab., 5 (1968), pp. 436–451.
- [39] G. F. NEWELL, *Queues with time-dependent arrival rates. II. The maximum queue and the return to equilibrium*, J. Appl. Probab., 5 (1968), pp. 579–590.
- [40] G. F. NEWELL, *Queues with time-dependent arrival rates. III. A mild rush hour*, J. Appl. Probab., 5 (1968), pp. 591–606.
- [41] G. F. NEWELL, *Applications of Queueing Theory*, Monogr. Appl. Probab. Statist., Chapman and Hall, London, 1971.
- [42] K. RAMANAN, *Reflected diffusions defined via the extended Skorokhod map*, Electron. J. Probab., 11 (2006), pp. 934–992.
- [43] M. H. ROTHKOPF AND S. S. OREN, *A closure approximation for the nonstationary $M/M/s$ queue*, Management Sci., 25 (1979), pp. 522–534.
- [44] T. SIMUNIC, L. BENINI, P. GLYNN, AND G. DE MICHELI, *Event-driven power management*, IEEE Trans. Computer-Aided Design, 20 (2001), pp. 840–857.
- [45] T. SIMUNIC, H. VIKALO, P. GLYNN, AND G. DE MICHELI, *Energy efficient design of portable wireless systems*, in Proceedings of the 2000 International Symposium on Low Power Electronics and Design, IEEE, Piscataway, NJ, 2000, pp. 49–54.
- [46] A. V. SKOROKHOD, *Stochastic equations for diffusion in a bounded region*, Theory Probab. Appl., 6 (1961), pp. 264–274.
- [47] W. WHITT, *Internet Supplement, Stochastic-Process Limits. An Introduction to Stochastic-Process Limits and Their Application to Queues*, Springer Ser. Oper. Res., <http://www.columbia.edu/~ww2040/supplement.html> (2002).
- [48] W. WHITT, *Stochastic-Process Limits. An Introduction to Stochastic-Process Limits and Their Application to Queues*, Springer Ser. Oper. Res., Springer-Verlag, New York, 2002.
- [49] R. J. WILLIAMS, *On dynamic scheduling of a parallel server system with complete resource pooling*, in Analysis of Communication Networks: Call Centres, Traffic and Performance (Toronto, ON, 1998), Fields Inst. Commun. 28, Amer. Math. Soc., Providence, RI, 2000, pp. 49–71.
- [50] J. YONG AND X. Y. ZHOU, *Stochastic Controls*, Appl. Math. 43, Springer-Verlag, New York, 1999.