

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 and, for a class of performance measures that satisfy a certain continuity property, it is then shown that any policy that is optimal for the fluid control problem is asymptotically optimal for the original network problem. Examples of performance measures with this property are presented, and simulations implementing proposed asymptotically optimal policies are presented. The use of directional derivatives of the reflection map for solving fluid optimal 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

1. Introduction. Most real-world queueing systems evolve according to laws that vary with time. The expository paper [27] outlines the applications of time-varying stochastic networks to telecommunications. In the context of computer engineering applications arise in the fields of power aware scheduling and temperature aware scheduling (see, e.g., [2, 38, 39]), as well as the design of web servers (see, e.g., [8]). For a broader range of applications pertaining to computer science, the reader is directed to [18] and references therein. Examples of other applications can be found in, e.g., [6, 17, 23, 37], while for work focusing on time-dependent phase-type distributions one should consult [31, 32] and references therein.

The focal point of the present paper is the rigorous study of certain aspects of 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 certain 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 (or FSLLN, see, e.g., Theorem 2.1 of [25]) limits of the original systems and lead to deterministic control problems. Second-order problems,

* Many thanks to 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.

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

[‡]Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA (kramanan@math.cmu.edu). This work was supported in part by the National Science Foundation under grants No. DMS-0405343 and CMMI-0728064. 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.

additionally, take into account certain fluctuations around the FSLLN limits. 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.

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. 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., [42] 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 pre-limit control problem (see [3, 4, 5] for examples). Other references in this context include [12, 14, 29, 30, 28] for the use of fluid control problems and [1, 19, 21, 24] 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 [25] is the cornerstone of the rigorous approach to the identification of both the first- and the second-order approximations for the $M_t/M_t/1$ queue under the *uniform acceleration* regime. The authors of [25] employ the theory of strong approximations (see, e.g., [9, 10, 16]) to develop a Taylor-like expansion of sample paths of queue lengths, establishing a FSLLN 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 [40] relaxes certain technical assumptions posited in [25] and exhibits some more general results. An off-shoot of the queue-length expansion developed in [25] is the study of the second-order approximation term as a *directional derivative* of the one-sided reflection map (in an appropriate topology on the path-space). With a view towards establishing analogous approximations for networks with time-inhomogeneous arrival and service rates, properties of directional derivatives of multi-dimensional reflection maps corresponding to a general class of queueing networks were established in [26]. The article [26] 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 [36], 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 [7], where the authors study an optimal resource allocation control problem for a (stochastic) fluid model with multiple classes, and a controller who dynamically schedules different classes in a system that experiences an 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 pre-limit 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 [4, 5]. In fact, even a concept akin to the notion of *fluid-scale asymptotic optimality* described in [28] 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 developing a suitable

methodology for asymptotic optimal control for time-inhomogeneous networks. In this paper, we consider controls that are arrival and/or service rates in time-varying single-class queueing networks. Our general goal is to determine if there is a systematic way to design high-performance controls for a time-varying network by carrying out an optimal control analysis of a related (fluid) approximation of the network. While this philosophy is similar to that used for time-homogeneous networks, the nature of the asymptotic approximation has to be modified so as to capture the time-varying behaviour. 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, converge to a deterministic fluid limit system in the strong sense. We refer to the systems in this sequence with uniformly accelerated rates as the *pre-limit systems*. With the view that optimal control problems for the fluid limit are typically more tractable than for the pre-limit, we wish to answer the following question:

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

The phrase “suffices to identify” above can be interpreted in many ways. For instance, one may resolve to use *exactly* a fluid-optimal discipline when controlling the pre-limit systems, one may try to formalize the fluid-optimal disciplines in terms of a *state dependent (feedback) rule* and then use this rule to control the pre-limit systems, or one may opt for a heuristic way to “tweak” fluid-optimal policies to perform well in the pre-limit. We choose to focus on the simplest of the above-mentioned options, i.e., we simply seek a characterization of the class of performance measures for which the fluid-optimal disciplines are also asymptotically optimal. This characterization is the main result of the present paper and is exhibited in Theorem 5.3.

While it is natural to expect that such a connection between the fluid and pre-limit optimal control problems exists, in Section 7.2 we describe several natural situations where this fails to hold. This underscores the need for a rigorous analysis to determine when this intuition is indeed valid. We also emphasize that the task of identifying a fluid-optimal policy is not always straightforward. One approach, exploiting the results of [26] on the directional derivatives of the oblique reflection map (ORM), is illustrated in Section 6.1.4. This calculus of variations type 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 asymptotically optimality is formulated in Section 3. The related fluid optimal control problem is described in Section 4. The question of characterization we posed earlier is formalized in Section 5 via the notion of *fluid-optimizability* of performance measures, and our main results are stated and proved. Section 6 is dedicated to examples of relevant fluid-optimizable performance measures including aggregate Lipschitz holding cost. Concluding remarks and, in particular, examples where the connection between the fluid and original control problems fails to hold, are given in Section 7. All auxiliary technical results are gathered in the Appendix.

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

- $\bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$;
- $\mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ denotes the set of all (a.s.-equivalence classes of) nonnegative random variables on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$;

- $\text{meas}(\cdot)$ denotes the Lebesgue measure on \mathbb{R} ;
- $\mathbb{L}^1[0, T]$ denotes the set of all integrable functions defined on $[0, T]$ (with respect to the Lebesgue measure);
- $\mathbb{L}_+^1[0, T]$ denotes the set of all non-negative 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

$$\mathcal{I}_t(f) = \left(\int_0^t f_1(s) ds, \dots, \int_0^t f_k(s) ds \right) \quad \text{for } t \in [0, T]$$

with $f = (f_1, \dots, f_k) \in (\mathbb{L}^1[0, T])^k$;

- \mathcal{D} 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 denotes the subset of \mathcal{D} containing all nondecreasing functions;
- $\mathcal{D}_{\uparrow, f}$ stands for the subset of \mathcal{D}_\uparrow 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} ;
- $\mathcal{B}(Y)$ denotes the Borel σ -algebra on the topological space Y .

For the sake of completeness, we provide the following definitions to be used in the sequel:

DEFINITION 1.1. *Let $R \in \mathbb{R}^{\kappa \times \kappa}$ have positive diagonal elements, and let x be in \mathcal{D}^κ . We say that a pair $(z, l) \in \mathcal{D}^\kappa \times \mathcal{D}_\uparrow^\kappa$ 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 $t \in [0, T]$,*

$$(i) \quad z(t) \geq 0;$$

$$(ii) \quad z(t) = x(t) + Rl(t);$$

$$(iii) \quad \int_0^t \mathbf{1}_{[z^i(s) > 0]} dl^i(s) = 0, \text{ for } i = 1, \dots, \kappa.$$

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

The existence and uniqueness of the ORM for a particular class of matrices R was proved in the seminal paper [20] for continuous functions. Those results can be directly extended to càdlàg functions to support the above definition (see, for example, Theorems 2.1 in [13]). Further discussion of the ORM can be found in [26].

REMARK 1.2. Depending on typographical convenience, we will use Z_t and $Z(t)$ (to denote the value of a process Z at time t) interchangeably throughout the text.

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 the sequel) 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 the model 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}) , where

- $\mathbf{E} = (E^{(i)}, i = 1, 2, \dots, \kappa) \in \mathcal{D}_{\uparrow, f}^\kappa$ stands for the vector of (cumulative) exogenous arrivals to each of the κ stations;

- $\mathbf{S} = (S^{(i,j)}, i = 1, 2, \dots, \kappa, j = 1, 2, \dots, \kappa + 1) \in \mathcal{D}_{\uparrow, f}^{\kappa^2 + \kappa}$ 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^{ij} 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 $j = \kappa + 1$, $S^{i, \kappa + 1}$ represents the (cumulative) number of jobs that would complete service at the 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 $\boldsymbol{\mu} = (\mu_1, \mu_2, \dots, \mu_\kappa) \in \mathbb{L}_+^1[0, T]^\kappa$, 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 [22]. 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 the 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 associated with the routing matrix P , in the sense of Remark 1.4 of [26], satisfies the [H-R] condition of Definition 1.2 of [26]. This assumption on P yields the well-definedness of the Oblique Reflection Problem (ORP) and Lipschitz continuity of the reflection mapping associated with the routing matrix P . For more details, the reader is directed to Theorem 3.1 of [26].

Specifically, 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

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

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

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

and the jobs that would leave the network as

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

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

REMARK 2.2.

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

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:

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

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;
- $\tilde{\mathbf{S}}_t = (\tilde{S}_t^{(k,j)}, 1 \leq k \leq \kappa, 1 \leq j \leq \kappa + 1) \in \mathcal{D}_{\uparrow, f}^{\kappa + \kappa^2}$ 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 $j = \kappa + 1$);
- $\mathbf{Z}_t = (Z_t^{(k)}, k = 1, 2, \dots, \kappa) \in \mathcal{D}^{\kappa}$ stands for the vector of queue-length processes in the κ stations.

It can be shown that the 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. Since the stochastic processes modeling the times of these events are PPPs, with probability one, there are at most a finite number of such events during the time interval $[0, T]$; 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 worthwhile to note that the above construction departs from the one common in 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 an 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

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

Standard arguments (see, e.g., [26]) can be used to show that \mathbf{Z} satisfies

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

where Γ^P denotes the multi-dimensional oblique reflection map 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 $J : \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ 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 $J : \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ that is bounded from below and Borel measurable with $\mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2}$ endowed with the Borel σ -algebra in the product M'_1 -topology is called a performance measure.*

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 [41].

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 modelling 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 *non-anticipating*, i.e., predictable with respect to the filtration $\{\mathcal{H}_t\}$ defined by

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

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 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 define \mathcal{A} to be 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 *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 look into deterministic (i.e., not state-dependent) controls.

For any $(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathcal{A}$, we 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

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

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

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

assuming the quantity above is well-defined. Concrete examples of such optimal control problems are provided in Sections 6.1 and 6.2.

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 so-called “fluid optimal control problem” introduced in Section 4). Moreover, as discussed in Section 6, 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

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

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 such that its 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

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

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

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

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 optimal 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

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

where

- I denotes the $\kappa \times \kappa$ -dimensional identity matrix;
- $Q = P^r$ stands for the transpose of the fixed routing matrix P ;
- $\mathcal{I}(\boldsymbol{\lambda}) \in (\mathcal{C}[0, T])^\kappa$ is the vector of mean exogenous arrivals to each of the κ stations;
- $\mathcal{I}(\boldsymbol{\mu}) \in (\mathcal{C}[0, T])^\kappa$ denotes the vector of mean *potential* service completions in the κ stations;
- $\bar{\mathbf{X}} \in (\mathcal{C}[0, T])^\kappa$ is the vector of mean netput processes in the κ stations;
- $\Gamma^P : \mathcal{D}_\uparrow^\kappa \rightarrow \mathcal{D}_\uparrow^\kappa$ is the *oblique reflection map* (ORM) generated by the *oblique reflection problem* (ORP) associated with the routing matrix P .

4.2. Fluid-limit performance measure. Given the definition of the sequence of performance measures J^n from (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

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

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 subset of $(\mathbb{L}_+^1[0, T])^{2\kappa}$ containing rates that respect any exogenous constraints that may be imposed.

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

Optimize the value of $\bar{J}(\boldsymbol{\lambda}, \boldsymbol{\mu})$ over the set \mathbb{A} .

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

DEFINITION 4.1. *A policy $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \in \mathbb{A}$ is said to be fluid optimal if $\bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \leq \bar{J}(\boldsymbol{\lambda}, \boldsymbol{\mu})$, for every $(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A}$.*

5. A Criterion for Identification of Asymptotically Optimal Controls.

The fluid optimal 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 all admissible sequences whose terms are identically equal to a fixed fluid-optimal policy be (average) asymptotically optimal?

Theorem 5.3 provides a sufficient condition for an affirmative answer to this question, which is formally phrased in terms of the following notion:

DEFINITION 5.1. *Let $J : \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ be a performance measure and let $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \in \mathbb{A}$ be fluid optimal for the associated fluid performance measure \bar{J} in the sense of Definition 4.1. If*

$$\liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] \geq 0, \quad \text{a.s.}, \quad \forall \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}, \quad (5.1)$$

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

$$\liminf_{n \rightarrow \infty} \mathbb{E}[J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] \geq 0, \quad \forall \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}. \quad (5.2)$$

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

ASSUMPTION 5.2. *The constraint set \mathbb{A} is bounded in the space $(\mathbb{L}_+^1[0, T])^{2\kappa}$.*

THEOREM 5.3. *Suppose the mapping $J : \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ is continuous with respect to the product M'_1 -topology on $\mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2}$. Then, a.s.,*

$$\lim_{n \rightarrow \infty} [\bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] = 0, \quad (5.3)$$

and J is fluid-optimizable.

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

REMARK 5.4. We can, in fact, immediately deduce a seemingly stronger result. Namely, if J has the continuity properties stated in Theorem 5.3, then the pre-limit

value functions converge, along with the performances of the fluid-optimal policies, to the fluid-optimal value. More precisely, let $V^n = \inf_{(\boldsymbol{\lambda}, \boldsymbol{\mu}) \in \mathbb{A}} J^n(\boldsymbol{\lambda}, \boldsymbol{\mu})$ be the value function associated with the n^{th} system. Then, clearly, for any fluid-optimal policy $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \in \mathbb{A}$, we have

$$\limsup_{n \rightarrow \infty} (V^n - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)) \leq 0.$$

On the other hand, if J has the continuity property stated in Theorem 5.3, then J is fluid-optimizable. Therefore, it follows from (5.1) that

$$\liminf_{n \rightarrow \infty} (V^n - \bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)) \geq 0.$$

The last two assertions, when combined with (5.3), show that

$$\lim_{n \rightarrow \infty} V^n = \bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*). \quad (5.4)$$

We now turn to the proof of Theorem 5.3.

Proof. First, recalling that the uniform topology is stronger than the M'_1 -topology (see, e.g., the beginning of Section 13.6.2 of [41]), and using the continuity in the M'_1 -topology of the mapping J , we conclude that J is continuous in the uniform topology.

Let $\{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}$ be an arbitrary admissible sequence. The left-hand side of the inequality in (5.1) can be expanded and bounded from below in the following fashion:

$$\begin{aligned} & \liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] \\ & \geq \liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] + \liminf_{n \rightarrow \infty} [\bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)]. \end{aligned} \quad (5.5)$$

Using (3.1) and (4.2), we can rewrite the last term in (5.5) in terms of J as

$$\begin{aligned} & \liminf_{n \rightarrow \infty} [\bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)] \\ & = \liminf_{n \rightarrow \infty} \left[J(\mathcal{I}(\boldsymbol{\lambda}^*), \text{diag}(\mathcal{I}(\boldsymbol{\mu}^*))\hat{P}) - J\left(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}^*), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}^*)\right) \right]. \end{aligned} \quad (5.6)$$

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

$$\left\| \frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}^*) - \mathcal{I}(\boldsymbol{\lambda}^*) \right\|_T \rightarrow 0, \quad \left\| \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}^*) - \text{diag}(\mathcal{I}(\boldsymbol{\mu}^*))\hat{P} \right\|_T \rightarrow 0. \quad (5.7)$$

More precisely, for every k and j , as $n \rightarrow \infty$, we have $\frac{1}{n}E^{(k)}(n\boldsymbol{\lambda}^*) \rightarrow \mathcal{I}(\lambda_k^*)$ and $\frac{1}{n}S^{(k,j)}(n\boldsymbol{\mu}^*) \rightarrow \mathcal{I}(p_{kj}\mu_k^*)$, \mathbb{P} -a.s., in the uniform topology on $\mathcal{D}([0, T])$. Due to (5.7) and the continuity of J in the uniform topology, the limit inferior in (5.6) is the proper limit and it is equal to zero. This establishes (5.3).

Let us now concentrate on the right-hand side of the inequality (5.5), and fix an $\omega \in \Omega$ for which (5.7) holds. All random quantities in the remainder of the proof will be evaluated at that ω without explicit mention. Due to fluid-optimality of $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$, this term dominates

$$\underline{j} := \liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)].$$

Without loss of generality, we can assume that $\underline{j} < \infty$. Let $\{(\eta_l, \nu_l)\}_{l \in \mathbb{N}}$ denote the subsequence of pairs $\{(\boldsymbol{\lambda}_{n_l}, \boldsymbol{\mu}_{n_l})\}_{l \in \mathbb{N}}$ along which the limit inferior above is attained as the proper limit.

Since Assumption 5.2 implies that $\{(\eta_i, \nu_i)\}_{i \in \mathbb{N}}$ are bounded in $(\mathbb{L}_+^1[0, T])^{2\kappa}$, by Lemma B.1 there exists a further subsequence $\{(\eta_i, \nu_i)\}_{i \in \mathbb{N}}$ and a function $F \in \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2}$ such that $(\mathcal{I}(\eta_i), \text{diag}(\mathcal{I}(\nu_i))\hat{P}) \rightarrow F$ in the product M_1' -topology on $\mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2}$. The assumed continuity of J in the product M_1' -topology and the definition of \bar{J} given in (4.2) yields

$$\lim_{i \rightarrow \infty} \bar{J}(\eta_i, \nu_i) = \lim_{i \rightarrow \infty} J(\mathcal{I}(\eta_i), \text{diag}(\mathcal{I}(\nu_i))\hat{P}) = J(F).$$

On the other hand, by Theorem A.1, the components of the random vector $(\frac{1}{n}\mathbf{E}(ne), \frac{1}{n}\mathbf{S}(ne))$ converge to the identity function e on $[0, T]$ in the uniform topology. So, we can utilize Lemma B.3 with the components of $(\frac{1}{n_i}\mathbf{E}(n_i e), \frac{1}{n_i}\mathbf{S}(n_i e))$ corresponding to the Y_n in the lemma and the components of (η_{n_i}, ν_i) corresponding to the ν_n in the lemma to conclude that

$$(\frac{1}{n_i}\mathbf{E}(n_i \eta_i), \frac{1}{n_i}\mathbf{S}(n_i \nu_i)) \rightarrow F \quad \text{as } i \rightarrow \infty,$$

in the product M_1' -topology. Hence, $\lim_{i \rightarrow \infty} J^{n_i}(\eta_i, \nu_i) = J(F)$. We conclude that $\bar{j} = 0$, which completes the proof of the first statement.

As for the average fluid-optimizability, note that due to the boundedness of the mapping J , the terms

$$J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J^n(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = J(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}_n), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}_n)) - J(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}^*), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}^*))$$

are bounded from below by a constant (say, by $-2L$, where L denotes the uniform upper bound on the mapping J). Hence, Fatou's lemma is applicable to the left-hand side of (5.2). This, along with the already proved inequality (5.1), completes the proof of the second statement. \square

If one is willing to impose a stricter - *uniform* - continuity condition in the above result, then one can relax the topology with respect to which continuity is required. To substantiate this claim, we need the following lemma, which is a direct consequence of Definition 3.1.

LEMMA 5.5. *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 such that*

$$\liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J_n^*] \geq 0, \quad \text{a.s., for all } \{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}} \in \mathcal{A}^{\mathbb{N}}, \quad (5.8)$$

and

$$\lim_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n^*, \boldsymbol{\mu}_n^*) - J_n^*] = 0, \quad \text{a.s.} \quad (5.9)$$

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.*

PROPOSITION 5.6. *Assume that J is uniformly continuous in the uniform topology on $\mathcal{D}_{\uparrow, f}^{2\kappa+\kappa^2}$. Then J is a fluid-optimizable performance measure.*

Moreover, if J is bounded and uniformly continuous in the uniform topology on $\mathcal{D}_{\uparrow, f}^{2\kappa+\kappa^2}$, then it is average fluid-optimizable.

Proof. We will reuse the notation established in the statement and proof of Theorem 5.3. Define the constant sequence $J_n^* = \bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ for every $n \in \mathbb{N}$. We have already shown (see expansion (5.6), the limit in (5.7) and the discussion following it) that the limit in (5.9) holds due to the continuity of the mapping J in the uniform topology. It remains to verify that condition (5.8) of Lemma 5.5 holds as well.

To see this, let us temporarily fix the admissible sequence $\{(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)\}_{n \in \mathbb{N}}$ and bound the term in (5.8) as follows:

$$\begin{aligned} & \liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - J_n^*] \\ & \geq \liminf_{n \rightarrow \infty} [J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)] + \liminf_{n \rightarrow \infty} [\bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)]. \end{aligned}$$

The second term on the right-hand side of the above display is a.s. nonnegative due to fluid-optimality of the policy $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$. As for the first term, we will prove that the limit inferior is a proper limit and equal to zero. Our tools are the Borel-Cantelli lemma and the submartingale inequality. Fix an arbitrary $\delta > 0$. Due to the uniform continuity of J , there exists a positive constant $\varepsilon(\delta)$ such that for every $x, y \in \mathcal{D}_{\uparrow, f}^{2\kappa+\kappa^2}$, $\|x - y\|_T < \varepsilon(\delta) \Rightarrow |J(x) - J(y)| < \delta$. Hence, for every n , we have

$$\begin{aligned} & \mathbb{P}[|J^n(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n) - \bar{J}(\boldsymbol{\lambda}_n, \boldsymbol{\mu}_n)| > \delta] \\ & = \mathbb{P}\left[\left|J\left(\frac{1}{n}\mathbf{E}(n\boldsymbol{\lambda}_n), \frac{1}{n}\mathbf{S}(n\boldsymbol{\mu}_n)\right) - J(\mathcal{I}(\boldsymbol{\lambda}_n), \text{diag}(\mathcal{I}(\boldsymbol{\mu}_n)) \cdot \hat{P})\right| > \delta\right] \\ & \leq \mathbb{P}\left[\left\|\frac{1}{n}(\mathbf{E}(n\boldsymbol{\lambda}_n), \mathbf{S}(n\boldsymbol{\mu}_n)) - (\mathcal{I}(\boldsymbol{\lambda}_n), \text{diag}(\mathcal{I}(\boldsymbol{\mu}_n)) \cdot \hat{P})\right\|_T > \varepsilon(\delta)\right]. \end{aligned} \quad (5.10)$$

Using Lemma A.2 and the expression for the fourth moment of a Poisson random variable, we can further bound the last expression in (5.10) 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, thanks to Assumption 5.2, we conclude that K_n are 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, the Borel-Cantelli lemma completes the proof of the first claim.

The same reasoning that was employed in the proof of Theorem 5.3 yields the second claim. \square

6. Applications of Fluid-optimizability Criteria. In this section, we illustrate how the concept of fluid-optimizability can be applied to study certain network optimal control problems. Specifically, we provide examples of optimal control problems for which the performance measure is fluid-optimizable and the fluid-optimal control policy can be explicitly determined. The latter is done with the use of directional derivatives of the multi-dimensional reflection map introduced in [26], a

technique that may be of independent interest. In addition, we discuss how the fluid-optimal policy can be used to design (near-optimal) controls for a given “pre-limit” system. In addition, we use simulations to explore the effect that the choice of the uniformly accelerated sequence into which the actual system of interest is embedded has on the resulting near-optimal control. The simulations also serve to illustrate the fluid-optimizability result of Theorem 5.3. In Section 6.1, we consider a holding cost performance measure, while in Section 6.2 we consider a variant that also incorporates throughput. In both cases, we assume that the open network has κ stations and that a 0–1 upper-triangular routing matrix P (i.e., it is feedforward and with deterministic routing).

6.1. An example with holding costs. In this section we consider a performance measure involving the 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 queue-lengths. In fact, cost structures similar to ours are quite standard (see, e.g., Chapter 7 of [19] or p. 60 of [43]). For more recent examples of similar cost functions, the reader is directed to [15] and references therein. It is worth noting that due to the fact that our time-horizon is finite, there is no discounting or time-averaging of the holding cost.

6.1.1. 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

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

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

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

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

LEMMA 6.1. *The mapping h defined in (6.1) is Lipschitz continuous on $\mathcal{D}_{\uparrow, f}^{2\kappa+\kappa^2}$ with respect to the uniform metric. If, in addition, Assumption 5.2 is satisfied, h is a fluid-optimizable performance measure.*

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

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

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

$$|h_k(Z^{(k)}(t)) - h_k(\tilde{Z}^{(k)}(t))| \leq C_k |Z^{(k)}(t) - \tilde{Z}^{(k)}(t)| \leq C_k \|Z^{(k)} - \tilde{Z}^{(k)}\|_T$$

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 [41]), 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 $\mathcal{D}_{\uparrow, f}^{2\kappa+\kappa^2}$.

In the presence of Assumption 5.2, we invoke Proposition 5.6 to conclude that the performance measure h is fluid-optimizable. \square

6.1.2. *The optimal control problem.* Consider a tandem queue with the processes E^A , S_1^A and S_2^A of exogenous arrivals to the first station and potential services at the first and second station, respectively, being modelled using PPPs and rates $\lambda^A, \mu_1^A, \mu_2^A \in \mathbb{L}_+^1[0, T]$, 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. The service rate μ_1^A in the first station serves as the control, while λ^A and μ_2^A are taken to be known (one can assume that λ^A and μ_2^A 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

$$J^A(E^A, S^A) = \int_0^T (h_1^A(Z_1^A(t)) + h_2^A(Z_2^A(t))) dt$$

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

- Z_i^A denotes the queue length of the i^{th} queue in the tandem for $i = 1, 2$, associated with the arrival and service processes (E^A, S_1^A, S_2^A) ;
- $h_1^A : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is given by $h_1^A(x) = c_1^A x^2$ for every $x \in \mathbb{R}_+$, with $c_1^A > 0$ constant;
- $h_2^A : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is given by $h_2^A(x) = c_2^A x$ for every $x \in \mathbb{R}_+$ and for a certain constant $0 < c_2^A < \mathcal{I}_T(\lambda^A)$.

It is natural to impose the following constraint on μ_1^A which ensures that admissible policies do not have more (mean) cumulative service available than there are (mean) cumulative arrivals:

$$\mathcal{I}_T(\mu_1^A) \leq \mathcal{I}_T(\lambda^A). \quad (6.3)$$

REMARK 6.2. The above set-up 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 h_1^A and h_2^A of the queue lengths are incurred. The controller's goal is to minimize the total holding cost J^A 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.

The superscript “A” used above refers to the fact that these quantities correspond to the actual network control problem of interest. Following the philosophy outlined in Section 5, in order to analyze this 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. Simulations illustrating the effect of the choice of the “embedding constant” N on the near-optimal control obtained are presented in Section 6.1.5. With an integer N that will serve as the embedding constant fixed, the construction of a uniformly accelerated sequence of systems described in Section 3 implies that in order for the actual system to correspond to the N^{th} system in the sequence, the “basic” arrival rate to the first station $\lambda \in \mathbb{L}_+^1[0, T]$ and the “basic” service rate μ_2 at the second station should be given by $\lambda = \frac{1}{N} \lambda^A$

and $\mu_2 = \frac{1}{N} \mu_2^A$. Moreover, the performance measure J for the “basic” system must take the form

$$J(E, S) = \int_0^T (h_1(Z_1(t)) + h_2(Z_2(t))) dt \quad \text{for } (E, S) \in (\mathcal{D}_{\uparrow, f}[0, T])^2 \quad (6.4)$$

with $h_i, i = 1, 2$, given by $h_1(x) = N^2 c_1^A x^2$ and $h_2(x) = N c_2^A x$ for every $x \in \mathbb{R}_+$, where Z_i denotes the queue length of the i^{th} queue in the tandem for $i = 1, 2$, associated with the triplet (E, S_1, S_2) of arrival and service processes (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 $J^A(E^A, S^A) = J^N(\lambda, \mu)$, 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 (6.3) pertaining to the actual system into the following constraint on the “basic” controls:

ASSUMPTION 6.3. *The constraint set is $\mathbb{A} = \{\mu \in \mathbb{L}_+^1[0, T] : \mathcal{I}_T(\mu) \leq \mathcal{I}_T(\lambda)\}$.*

6.1.3. *The related fluid optimal control problem.* Since h is fluid-optimizable, it follows from Definitions 3.1 and 4.1 that to identify a strongly asymptotically optimal sequence for a control problem with h as performance measure, it suffices to analyze the corresponding fluid optimal control problem. We illustrate this procedure for the control problem introduced in Section 6.1.2, using calculus of variations type techniques.

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 $\mu_2 \in \mathbb{L}_+^1[0, T]$. Assume that there are no exogenous arrivals to the second station. Our fluid optimal 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 (6.4), we define the fluid-limit holding cost as

$$\bar{h}(\mu) = \int_0^T [h_1(\bar{Z}_t^1(\mu)) + h_2(\bar{Z}_t^2(\mu))] dt \quad \text{for every } \mu \in \mathbb{A}, \quad (6.5)$$

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 $c_1 = N^2 c_1^A$ and $c_2 = N c_2^A$ to simplify the notation, while $\bar{Z}^i(\mu), i = 1, 2$ denote the queue lengths in the fluid tandem queue (as a function of μ).

6.1.4. *Solution of the fluid optimal control problem.* In the present section, we identify a fluid-optimal control. To keep the calculations as simple as possible and make the illustration of our calculus of variations-like approach to the problem transparent, we additionally set $\mu_2^A \equiv 0$. The explicit form of the directional derivative of the ORM obtained in [26] 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 tradeoff between the marginal costs in the two stations to be considered. Heuristically, one needs to identify 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. As we are going to see shortly, this intuition coincides with the formal solution.

LEMMA 6.4. *The policy $\hat{\mu} \in \mathbb{A}$, defined by*

$$\hat{\mu} = \lambda \mathbf{1}_{[t_c, T]} \quad \text{with } t_c = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) \geq c := \frac{c_2}{2c_1}\}$$

is fluid optimal for the fluid optimal control problem of Section 6.1.3.

Proof. Suppose that a fluid-optimal policy exists and denote it by $\hat{\mu}$. We shall first argue that the following claim holds:

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

Proof of Claim 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 is well known that the queue-length in the first station, when $\hat{\mu}$ is the service employed there, can be rewritten as

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

while the queue-length in the second station equals

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

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

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

Then, $\mathcal{I}_t(\lambda) - \mathcal{I}_t(\check{\mu}) = \bar{Z}_t^1(\hat{\mu}) \geq 0$ for every t . Moreover, $\bar{Z}_t^i(\hat{\mu}) = \bar{Z}_t^i(\check{\mu})$ for $i = 1, 2$ and $t \in [0, T]$. Hence, $h(\hat{\mu}) = h(\check{\mu})$ while $\check{\mu}$ satisfies the desired inequality.

Let us return to the proof of the lemma assuming that $\hat{\mu}$ satisfies the inequality in Claim 1. If $\hat{\mu}$ is fluid optimal, then for every perturbation $\Delta\mu \in \mathbb{L}^1[0, T]$ and for every constant $\varepsilon > 0$ such that $\hat{\mu} + \varepsilon\Delta\mu \in \mathbb{A}$, we must have

$$\bar{h}(\hat{\mu} + \varepsilon\Delta\mu) - \bar{h}(\hat{\mu}) \geq 0. \quad (6.6)$$

From the above equations for \bar{Z} , it is clear that for any $\mu \in \mathbb{A}$ that satisfies the condition of Claim 1,

$$\bar{Z}(\mu) = (\bar{Z}^1(\mu), \bar{Z}^2(\mu)) = \Gamma(\bar{X}(\mu)) = \Gamma(\mathcal{I}(\lambda - \mu), \mathcal{I}(\mu - \mu_2)). \quad (6.7)$$

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

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

where, as in [26], 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} 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 (6.6) holds if and only if

$$\begin{aligned} \frac{1}{\varepsilon} \int_0^T [c_1(\bar{Z}_t^1(\hat{\mu} + \varepsilon\Delta\mu) - \bar{Z}_t^1(\hat{\mu}))(\bar{Z}_t^1(\hat{\mu} + \varepsilon\Delta\mu) + \bar{Z}_t^1(\hat{\mu})) \\ + c_2(\bar{Z}_t^2(\hat{\mu} + \varepsilon\Delta\mu) - \bar{Z}_t^2(\hat{\mu}))] dt \geq 0. \end{aligned} \quad (6.8)$$

It follows from Theorem 1.6.2 in [26] that, as $\varepsilon \downarrow 0$, the pointwise limit of $\nabla_\chi^\varepsilon \Gamma(\bar{X}(\hat{\mu}))$ exists and is given explicitly by

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

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 $\Phi^i(t) \doteq \{s \leq t : \bar{Z}_s^i(\hat{\mu}) = 0\} = \{s \leq t : \bar{X}_s^i(\hat{\mu}) = 0\}$, $i = 1, 2$, 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 fact that both queues in the tandem are non-idling (see *Claim 1*).

Therefore, for every perturbation $\Delta\mu$ such that $\mu + \varepsilon\Delta\mu \in \mathbb{A}$ for all sufficiently small $\varepsilon > 0$ (we shall refer to such perturbations as admissible perturbations), taking limits as $\varepsilon \downarrow 0$ in (6.8), we see that

$$\int_0^T [(\bar{Z}_t^1(\hat{\mu}) - c)(-\mathcal{I}_t(\Delta\mu) + \gamma^1(t)) + c\gamma^2(t)] dt \geq 0, \quad (6.9)$$

with $c := \frac{c_2}{2c_1}$. Define the time instances

$$\hat{t} = \inf\{t \in [0, T] : \mathcal{I}_t(\hat{\mu}) > 0\} \quad \text{and} \quad t_c = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) \geq c\}.$$

Due to the assumption that $\mu_2 \equiv 0$, we immediately conclude that $\Phi^2(t) = [0, t \wedge \hat{t}]$ for every $t \in [0, T]$. We now claim that:

Claim 2: $t_c \leq \hat{t}$.

Proof of Claim 2: Consider a perturbation $\Delta\mu$ such that $\Delta\mu(t) \leq 0$ for all $t \in [0, T]$ and $\Delta\mu(t) = 0$ for every $t \in [0, \hat{t}]$. Then, $\Delta\mu$ is clearly an admissible perturbation and, moreover, $\gamma^1 \equiv 0$ and $\gamma^2(t) = \sup_{s \in [0, t \wedge \hat{t}]} [-\mathcal{I}_s(\Delta\mu)] = 0$. Therefore, (6.9) reduces to

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

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

We now show that, in fact:

Claim 3: $t_c = \hat{t}$.

Proof of Claim 3: Let us assume that $t_c < \hat{t}$ 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 (6.9) can be rewritten as

$$\int_0^T [(\bar{Z}_t^1(\hat{\mu}) - c)(-\mathcal{I}_t(\Delta\mu) + \mathcal{I}_{m_1(t)}(\Delta\mu))] dt \geq 0, \quad \text{for every } \Delta\mu \geq 0.$$

Thus,

$$\int_0^T \int_0^T [(\bar{Z}_t^1(\hat{\mu}) - c)(-\mathbf{1}_{[m_1(t), t]}(u))\Delta\mu(u)] du dt \geq 0, \quad \text{for every } \Delta\mu \geq 0.$$

Due to Fubini's theorem, the above inequality yields

$$\int_0^T \Delta\mu(u) \left(\int_0^T [(\bar{Z}_t^1(\hat{\mu}) - c)(-\mathbf{1}_{[m_1(t), t]}(u))] dt \right) du \geq 0, \quad \text{for every } \Delta\mu \geq 0.$$

We define the function $F : [0, T] \rightarrow \mathbb{R}_+$

$$F(u) = \int_0^T [(\bar{Z}_t^1(\hat{\mu}) - c)(-\mathbf{1}_{[m_1(t), t]}(u))] dt$$

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

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

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

To conclude the proof of the lemma, it suffices to show

Claim 4: $\hat{\mu}(t) = \lambda(t)$ for almost every $t \geq \hat{t}$.

Proof of Claim 4: Let us assume that there exists a pair of time instances $t_1 < t_2$ such that $\hat{t} < t_1$ and $\bar{Z}_t^1(\hat{\mu}) > c$ for every $t \in (t_1, t_2)$. Note that $\Phi^1(t) \subset (t_1, t_2)^c$ and recall that $\Phi^2(t) = [0, t \wedge \hat{t}]$ 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 (6.9) reads as

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

From the choice of $\Delta\mu$ and the definition of t_1 and t_2 , we conclude that and the above expression must be strictly negative, which contradicts the inequality (6.9). Thus, $\bar{Z}_t^1(\hat{\mu}) \leq c$ for every $t \in (\hat{t}, T)$.

Using an analogous argument, one can show that $\bar{Z}_t^1(\hat{\mu}) \geq c$ for every $t \in (\hat{t}, T)$. Indeed, assume the contrary and set

$$\begin{aligned} t_1 &= \inf\{t > \hat{t} : \bar{Z}_t^1(\hat{\mu}) < c\} \wedge T, \\ t_2 &= \inf\{t > t_1 : \bar{Z}_t^1(\hat{\mu}) = c\} \wedge \inf\{t > t_1 : \bar{X}_t^1(\hat{\mu}) = 0\} \wedge T. \end{aligned}$$

With this choice of t_1 and t_2 , we again have $\Phi^1(t) \subset (t_1, t_2)^c$. Let $\Delta\mu$ be an admissible permutation such that $\Delta\mu(t) = 0$ for $t \in (t_1, t_2)$, $\mathcal{I}_{t_1}(\Delta\mu) = \mathcal{I}_{t_2}(\Delta\mu) = 0$ and $\mathcal{I}_t(\Delta\mu) < 0$ for $t \in (t_1, t_2)$. This choice of $\Delta\mu$ implies $\gamma^1(t) = \gamma^2(t) = 0$ for all t , and the left hand side of (6.9) becomes

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

Using the negativity of $\mathcal{I}_t(\Delta\mu)$ for $t \in (t_1, t_2)$ and the definition of t_1 and t_2 , we conclude that the above expression is negative which contradicts (6.9). Hence, $\bar{Z}_t^1(\hat{\mu}) \geq c$ for $t \in (\hat{t}, T)$.

Combining the above two inequalities, we conclude that $\bar{Z}_t^1(\hat{\mu}) = c$ for $t \in (\hat{t}, T)$. So, $\hat{\mu}(t) = \lambda(t)$, for almost every $t \in (\hat{t}, T)$. This proves the fourth claim and, thus, concludes the proof of the lemma. \square

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

COROLLARY 6.5. *The fluid-optimal value for the fluid optimal control problem of Section 6.1.3 is given by*

$$\bar{h}(\hat{\mu}) = c_1^A \int_0^{t_c^A} (\mathcal{I}_t(\lambda^A))^2 dt + c_2^A \int_{t_c^A}^T (\mathcal{I}_t(\lambda^A) - \frac{c^A}{2}) dt$$

with $t_c^A = \inf\{t > 0 : \mathcal{I}_t(\lambda^A) > c^A := \frac{c_2^A}{2c_1^A}\}$ and c_1^A, c_2^A and λ^A as in Section 6.1.2.

Proof. Using the form of the fluid-optimal policy $\hat{\mu}$ obtained in Lemma 6.4, we have that for every $t \in [0, T]$,

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

Hence,

$$\begin{aligned} \bar{h}(\hat{\mu}) &= \int_0^{t_c} c_1 (\mathcal{I}_t(\lambda))^2 dt + \int_{t_c}^T [c_1 c^2 + c_2 (\mathcal{I}_t(\lambda) - c)] dt \\ &= c_1 \int_0^{t_c} (\mathcal{I}_t(\lambda))^2 dt + c_2 \int_{t_c}^T (\mathcal{I}_t(\lambda) - \frac{c}{2}) dt. \end{aligned} \tag{6.10}$$

Recalling that $\lambda = \frac{1}{N} \lambda^A$, $c_1 = N^2 c_1^A$ and $c_2 = N c_2^A$, we get that

$$\begin{aligned} c &= \frac{N c_2^A}{2N^2 c_1^A} = \frac{c^A}{N}, \\ t_c &= \inf\{t > 0 : \mathcal{I}_t(\lambda) > c\} = \inf\{t > 0 : \mathcal{I}_t(\lambda^A) > c^A\} = t_c^A. \end{aligned}$$

With this in mind, the expression for the fluid-optimal cost of (6.10) becomes

$$\begin{aligned} \bar{h}(\hat{\mu}) &= N^2 c_1^A \int_0^{t_c^A} \frac{1}{N^2} (\mathcal{I}_t(\lambda^A))^2 dt + N c_2^A \int_{t_c^A}^T \frac{1}{N} (\mathcal{I}_t(\lambda^A) - \frac{c^A}{2}) dt \\ &= c_1^A \int_0^{t_c^A} (\mathcal{I}_t(\lambda^A))^2 dt + c_2^A \int_{t_c^A}^T (\mathcal{I}_t(\lambda^A) - \frac{c^A}{2}) dt. \end{aligned} \quad \square$$

6.1.5. Implementation. Lemmas 6.1 and 6.4, Assumption 6.3 and Proposition 5.6, 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 6.1.2. The details of the embedding procedure preceding the formulation and the solution of the fluid optimal control problem are described in Sections 6.1.2 and 6.1.4, respectively.

To illustrate the performance of the fluid-optimal discipline we obtained above, we ran simulations of the pre-limit systems when the fluid-optimal policy is implemented. All the simulations were conducted in C++ and the graphs were produced by R. We set a time horizon at $T = 1$ and conducted the simulations for the periodic arrival rate $\lambda^A(t) = 100(1 + \sin(10t))$, 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 $c_1^A = 1/20,000$ and $c_2^A = 1/200$. 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 (6.5) with constants $c_1 = 2500c_1^A = 0.125$ and $c_2 = 50c_2^A = 0.25$. Using Lemma 6.4, we obtain a fluid-optimal control $\hat{\mu} = \lambda \mathbf{1}_{[t_c, T]}$

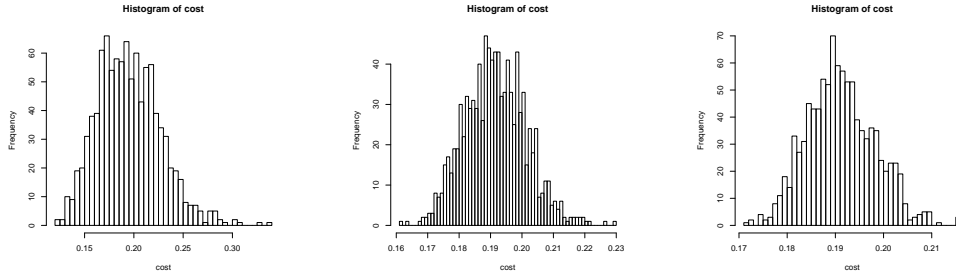


FIG. 6.1. Histograms of costs realized for the embedding constant of 50, $\lambda^A(t) = 100(1 + \sin(10t))$, $t \in [0, 1]$, and for uniform acceleration coefficients $n = 50, 500, 1000$. An approximate fluid-optimal value is 0.190846 (calculated using Mathematica).

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

TABLE 6.1

Summary statistics for 1000 simulations of the cost in case of the embedding constant of 50, $\lambda^A(t) = 100(1 + \sin(10t))$ and for acceleration coefficients in the first column.

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 6.1 and Table 6.1 summarize the results of applying the fluid-optimal policy $\hat{\mu}$ to the pre-limit systems. The embedding index of 50 is the one corresponding to the actual system in the sense of 6.1.2 and the outcome of the simulations of the cost of applying the fluid-limit optimal policy to the actual system can be seen in the leftmost graph in Figure 6.1.

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 $\lambda^A \equiv 100$, or the periodic arrival rate $\lambda^A(t) = 100(1 + \sin(10t))$, $t \in [0, T]$. The fluid performance measure \bar{h} is again as in (6.5), but now with constants $c_1 = c_2 = 0.5$. According to Lemma 6.4, the fluid-optimal policy in this case 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 6.2 and 6.3, and the summary statistics are provided in Tables 6.2 and 6.3. The simulated costs of employing the fluid-optimal policy in the actual system are given in the middle graphs on Figures 6.2 and 6.3. The reader interested in comparing the effects of different embedding constants should compare the left-most graph in Figure 6.1 to the middle graph in Figure 6.3. Figures 6.4 and 6.5 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 fluid-optimal service begins in the first station and starts “matching” the arrivals to the first station.

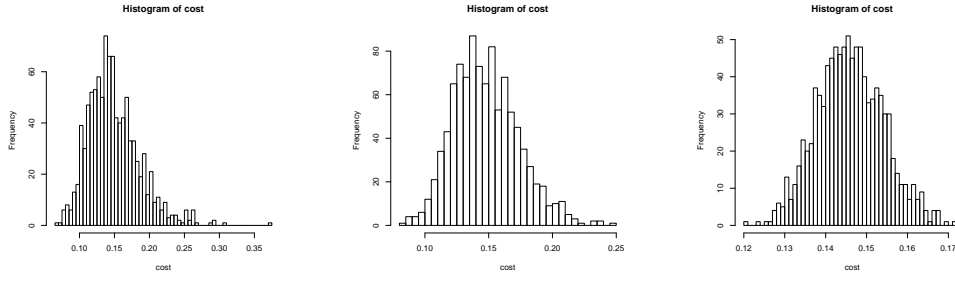


FIG. 6.2. Histograms of costs realized for the embedding constant of 100, $\lambda^A \equiv 100$ and for uniform acceleration coefficients $n = 50, 100, 1000$. The fluid-optimal value is $7/48 \approx 0.14583$.

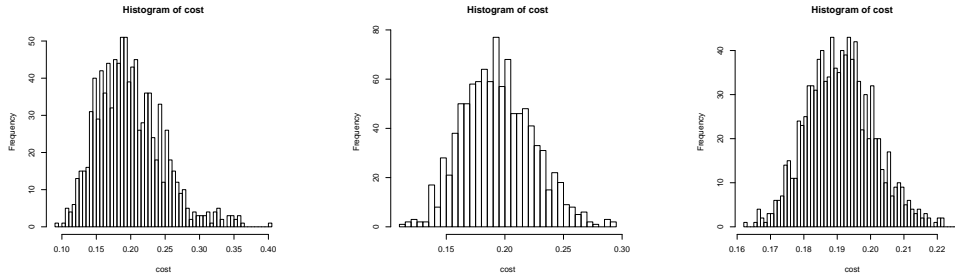


FIG. 6.3. Histograms of costs realized for the embedding constant of 100, $\lambda^A(t) = 100(1 + \sin(10t))$, $t \in [0, 1]$, and for uniform acceleration coefficients $n = 50, 100, 1000$. An approximate fluid-optimal value is 0.190846 (calculated using Mathematica).

REMARK 6.6. Note that the simulation results in the present section indeed illustrate the claim of Proposition 5.6 and Remark 5.4. In particular, the simulation values become more concentrated around their averages which, in turn, approach the theoretical fluid-optimal value, which also equals the limit of the pre-limit value functions.

6.2. Trade-off between holding cost and throughput. In this section, we consider a variation on the optimal control problem from Section 6.1 which, in addition to the holding cost, takes into account a reward for the completion of jobs during the interval $[0, T]$. The controller's goal is to balance the holding cost penalty with the profit generated by the completed jobs.

This can be viewed as a model of inventory control, in a setting similar to that described in Remark 6.2, except that this time the controller is in charge of a single station with a holding cost which is an increasing function of the number of jobs in the queue; on the other hand, there is revenue for all products that get out of the station which offsets the holding cost.

6.2.1. The performance measure and the optimal control problem. The aggregate holding cost h associated with the pair (\mathbf{E}, \mathbf{S}) was defined in (6.1). Let the profit generated by the completion of jobs during the time interval $[0, T]$ be given by a Lipschitz continuous function $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

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

TABLE 6.2

Summary statistics for 1000 simulations of the cost in case of the embedding constant of 100, $\lambda^A \equiv 100$ and for acceleration coefficients 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

TABLE 6.3

Summary statistics for 1000 simulations of the cost in case of the embedding constant of 100, $\lambda^A(t) = 100(1 + \sin(10t))$ and for acceleration coefficients in the first column.

We introduce a performance measure $J : \mathcal{D}_{\uparrow, f}^{2\kappa + \kappa^2} \rightarrow \mathbb{R}$ as

$$J(\mathbf{E}, \mathbf{S}) = h(\mathbf{E}, \mathbf{S}) - p\left(\sum_{k=1}^{\kappa} E^k(T) - \sum_{k=1}^{\kappa} Z^k(T)\right).$$

Due to the Lipschitz continuity of both the mapping p and the one-sided reflection map, one can use the same rationale used in the proof of fluid-optimizability of h to verify the uniform continuity of J . Proposition 5.6 then shows that the performance measure J is fluid-optimizable whenever the set of admissible controls is bounded in $(\mathbb{L}_+^1[0, T])^{2\kappa}$. Similarly to the previous example, the validity of Assumption 5.2 will be enforced additionally for the particular optimal control problem we look at next.

Let us consider a single station with a given service rate $\mu \in \mathbb{L}_+^1[0, T]$. 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 λ in the first station. In the proposed application above, it is natural to assume that the cumulative mean arrivals of materials into a production station do not greatly exceed the available cumulative service, and so we define the constraint set as $\mathbb{A} = \{\lambda \in \mathbb{L}_+^1[0, T] : \mathcal{I}_T(\lambda) \leq 2\mathcal{I}_T(\mu)\}$.

6.2.2. A related fluid optimal control problem and its solution. 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). The fluid optimal control problem consists of minimizing \bar{J} across $\lambda \in \mathbb{A}$.

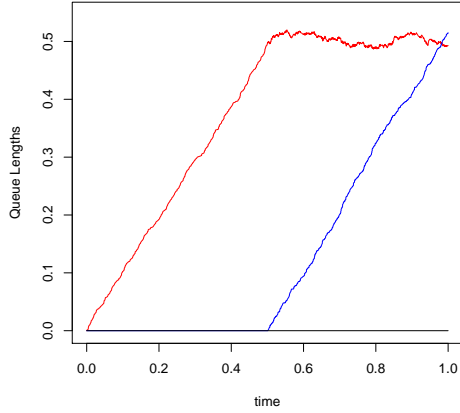


FIG. 6.4. One trajectory of the queue lengths in the first (increasing in the beginning) and second (the other curve) stations for the embedding constant of 100, the uniform acceleration coefficient $n = 1000$ and $\lambda^A \equiv 100$. The time at which service in the first station begins is 0.5.

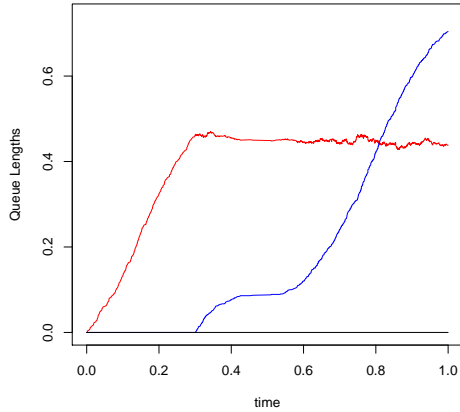


FIG. 6.5. One trajectory of the queue lengths in the first (increasing in the beginning) and second (the other curve) stations for the embedding constant of 100, the uniform acceleration coefficient $n = 1000$ and $\lambda^A \equiv 100(1 + \sin(10t))$. The time at which service in the first station begins is approximately 0.3.

LEMMA 6.7. *The policy $\hat{\lambda} = \mu$ is fluid optimal for the above fluid optimal control problem.*

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

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

The policy $\hat{\lambda} = \mu$ attains this lower bound and is, hence, fluid optimal. \square

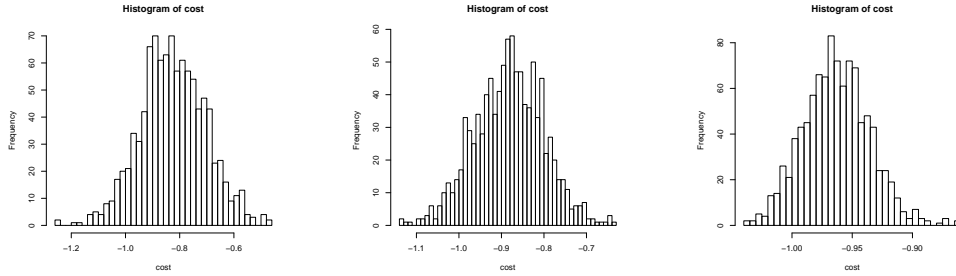


FIG. 6.6. Histograms of costs realized for the “basic” service rate $\mu \equiv 1$ and for uniform acceleration coefficients $n = 50, 100, 1000$. The fluid-optimal value equals -1 .

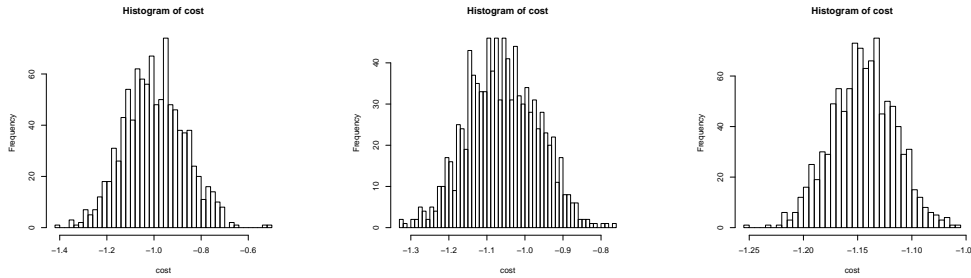


FIG. 6.7. Histograms of costs realized for the basic service rate $\mu(t) = 1 + \sin(10t)$ and for uniform acceleration coefficients $n = 50, 100, 1000$. The fluid-optimal value is $1.1 - \cos(10) \approx -1.184$.

6.2.3. Implementation. As an illustration of the performance of the fluid-optimal discipline we obtained in Lemma 6.7, we ran simulations of the pre-limit 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 case of the constant service rate are depicted in Figure 6.6 and in the case of periodic μ in Figure 6.7. The summary statistics are collected in Tables 6.4 and 6.5, for constant and periodic service rates, respectively.

REMARK 6.8. The approach of the simulated values to the theoretical limiting cost is slower than in the previous example. So, we included the results of taking a large uniform acceleration coefficient of 10,000 (see Figure 6.8). We believe that this is due to the effect of the system being continuously in heavy-traffic (under the fluid-optimal discipline). 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} .

7. Concluding remarks and further research. In this section, we note some features we encountered in this work which are unique to the time-inhomogeneous

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 6.4

Summary statistics for 1000 simulations of the cost in case of $\mu \equiv 1$ and for acceleration coefficients 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

TABLE 6.5

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

set-up. Some of these issues hint at possible directions of future research. Also, we broadly outline particular problems which were among the topics of [11].

7.1. 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 systems in heavy traffic; see, e.g., [42] 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, fluid-optimal policies 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 Poisson point processes (PPPs) proved to be a convenient modelling 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

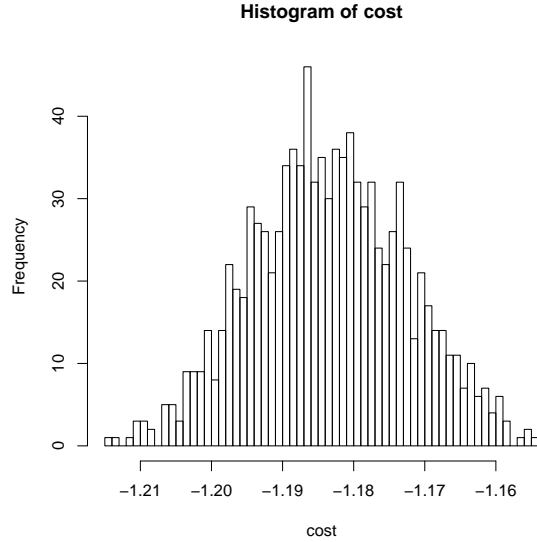


FIG. 6.8. Histogram of costs in the optimal control problem aimed at balancing the holding cost and the throughput, realized for $\mu(t) = 1 + \sin(10t)$ and for the uniform acceleration coefficient $n = 10,000$.

(see, e.g., Section 5.5 of [41] 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 pre-limit sequence of uniformly accelerated systems should be. The question of 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 6 (see, also, Corollary 6.5), 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 pre-limit stochastic optimal control problems and compare the performance of the fluid-optimal policies to the performance of the optimal control for the actual model.

7.2. Pertinent examples in earlier work. It may be intuitive to expect that fluid-optimal policies would provide near-optimal policies for some performance measures, and indeed such heuristics are employed by practitioners (see, e.g., [33, 34, 35]). 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. In [11], 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 fluid-optimal policies 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 fluid-optimality, as well as a broad class of asymptotically optimal sequences of policies, were identified (see Theorem 3.2.5 (p.40) and Theorem 3.4.8 (p.49), respectively, in [11]). Using these results, it is easy to show that for a certain set of parameters most, but **not all**, fluid-optimal policies are asymptotically optimal. In addition, it is also possible to construct an example (not studied in [11]) for which there is a unique fluid-optimal policy 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 set-up, 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 (p.91) of [11]). 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 (FSLLN).

In this section, we present and prove a version of the Functional Strong Law of Large Numbers (FSLLN). We emphasize that this result, albeit very similar in spirit to Theorem 2.1 of [25], is different. Stochastic processes used to model the exogenous arrival and potential service processes in [25] 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 $\mu \in \mathbb{L}_+^1[0, T]$ and $p : [0, T] \rightarrow [0, 1]$ be deterministic measurable functions 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. Let the sequence of stochastic processes $\{Y^{(n)}\}$ be defined as*

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

Then, as $n \rightarrow \infty$,

$$\frac{1}{n}Y^{(n)} \rightarrow \mathcal{I}((1-p)\mu), \text{ a.s. in the uniform topology.} \quad (\text{A.1})$$

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 Poisson point process 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 Poisson point process 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 Poisson point process associated with the process Y . By the Uniqueness Theorem for Laplace transforms and Lemma 12.1. in [22], 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 [22], we have that for every nonnegative, Borel measurable $f : \mathcal{S} \times \{0, 1\} \rightarrow \mathbb{R}_+$

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

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

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

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 [22] again to obtain

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

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

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

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

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

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

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

Let us define the operator F on real functions on \mathcal{S} to real functions on $\mathcal{S} \times \mathcal{T}$ 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 $g : \mathcal{S} \rightarrow \mathbb{R}_+$

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

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

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

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

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

2° Let g be a simple function of the form $g = \sum_{m=1}^M c_m \mathbf{1}_{B_m}$, 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

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

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 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

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

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 Poisson point process $\hat{\chi}$ acts on nonnegative measurable functions $g : [0, T] \rightarrow \mathbb{R}_+$ in the following way:

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

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

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

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

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

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

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

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

We continue with an application of the submartingale inequality.

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

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

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

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

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 [41], 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}^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 $\{\mathcal{I}_{q_1}(f_{n_l}^1)\}_l$ that converges in \mathbb{R} . The sequence $\{\mathcal{I}_{q_1}(f_{n_l}^2)\}_l$ 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 $\{G_l\}$ 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

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

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

$$F(q) = \lim_{l \rightarrow \infty} G_l(q).$$

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 $\{G_l\}$ 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 $\{G_l\}$ 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 [41]. So, we conclude that

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

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

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

Since by Theorem 12.5.2 of [41], the weak M_1 -topology coincides with the product M_1 -topology on \mathcal{D}^k , it suffices to again utilize Theorem 13.6.3. of [41] to conclude that the components of $\{G_l\}$ 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*

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

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 [41] 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 non-compact domains from p. 414 of [41]. In summary, it suffices to verify that

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

for ε that are positive continuity points of ν . The last claim is a direct consequence of Theorem 13.2.4 in [41], which completes the proof. \square

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] N. BANSAL, T. KIMBREL, AND K. PRUHS, *Dynamic speed scaling to manage energy and temperature*, in FOCS 2004, 2004, pp. 520–529.
- [3] 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.
- [4] 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.
- [5] A. BUDHIRAJA, A. P. GHOSH, AND C. LEE, *An ergodic rate control problem for single class queueing networks*. preprint, 2008.
- [6] M. J. CARRILLO, *Extensions of Palm’s theorem: A review*, Management Science, 37 (1991), pp. 739–744.
- [7] 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.
- [8] 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*, Journal of Engineering Application of Artificial Intelligence, 17 (2004), pp. 841–854.
- [9] M. CSÖRGŐ AND L. HORVÁTH, *Weighted approximations in probability and statistics*, Wiley Series in Probability and Mathematical Statistics: Probability and Mathematical Statistics, John Wiley & Sons Ltd., Chichester, 1993. With a foreword by David Kendall.
- [10] M. CSÖRGŐ AND P. RÉVÉSZ, *Strong approximations in probability and statistics*, Probability and Mathematical Statistics, Academic Press Inc. [Harcourt Brace Jovanovich Publishers], New York, 1981.
- [11] M. ČUDINA, *Asymptotically Optimal Control for Some Time-varying Stochastic Networks*, PhD thesis, Carnegie Mellon University, 2006. <http://www.math.utexas.edu/users/mcudina/thesis.pdf>.
- [12] P. DUPUIS, *Explicit solution to a robust queueing control problem*, SIAM J. Control Optim., 42 (2003), pp. 1854–1875 (electronic).
- [13] P. DUPUIS AND K. RAMANAN, *Convex duality and the Skorokhod problem. I, II*, Probab. Theory Related Fields, 115 (1999), pp. 153–195, 197–236.
- [14] ———, *An explicit formula for the solution of certain optimal control problems on domains with corners*, Teor. Ľmovir. Mat. Stat., (2000), pp. 32–48.
- [15] J. M. GEORGE AND J. M. HARRISON, *Dynamic control of a queue with adjustable service rate*, Oper. Res., 49 (2001), pp. 720–731.
- [16] 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.
- [17] L. V. GREEN AND P. J. KOLESAR, *A note on approximating peak congestion in $m_t/g/\infty$ queues with sinusoidal arrivals*, Man. Sci., 44 (1998), pp. 137–144.
- [18] 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, New York, NY, USA, 2006, ACM, pp. 203–215.
- [19] J. M. HARRISON, *Brownian motion and stochastic flow systems*, Robert E. Krieger Publishing Co. Inc., Malabar, FL, 1990. Reprint of the 1985 original.
- [20] J. M. HARRISON AND M. I. REIMAN, *Reflected Brownian motion on an orthant*, Ann. Probab., 9 (1981), pp. 302–308.
- [21] 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.
- [22] O. KALLENBERG, *Foundations of modern probability*, Probability and its Applications (New York), Springer-Verlag, New York, second ed., 2002.
- [23] 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.
- [24] H. J. KUSHNER, *Heavy traffic analysis of controlled queueing and communication networks*, vol. 47 of Applications of Mathematics (New York), Springer-Verlag, New York, 2001. Stochastic Modelling and Applied Probability.
- [25] A. MANDELBAUM AND W. A. MASSEY, *Strong approximations for time-dependent queues*, Math. Oper. Res., 20 (1995), pp. 33–64.
- [26] A. MANDELBAUM AND K. RAMANAN, *Directional derivatives of oblique reflection maps*, Mathematics of Operations Research (to appear), (2006).

- [27] W. A. MASSEY, *The analysis of queues with time-varying rates for telecommunications models*, Telecommunication Systems, 21 (2002), pp. 173–204.
- [28] S. MEYN, *Control techniques for complex networks*, Cambridge University Press, Cambridge, 2008.
- [29] S. P. MEYN, *Sequencing and routing in multiclass queueing networks. I. Feedback regulation*, SIAM J. Control Optim., 40 (2001), pp. 741–776 (electronic).
- [30] ———, *Sequencing and routing in multiclass queueing networks. II. Workload relaxations*, SIAM J. Control Optim., 42 (2003), pp. 178–217 (electronic).
- [31] 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.
- [32] ———, *The $[Ph_t/Ph_t/\infty]^K$ queueing system. II. The multiclass network*, INFORMS J. Comput., 16 (2004), pp. 275–283.
- [33] G. F. NEWELL, *Queues with time-dependent arrival rates. I. The transition through saturation*, J. Appl. Probability, 5 (1968), pp. 436–451.
- [34] ———, *Queues with time-dependent arrival rates. II. The maximum queue and the return to equilibrium*, J. Appl. Probability, 5 (1968), pp. 579–590.
- [35] ———, *Queues with time-dependent arrival rates. III. A mild rush hour*, J. Appl. Probability, 5 (1968), pp. 591–606.
- [36] ———, *Applications of queueing theory*, Chapman and Hall Ltd., London, 1971. Monographs on Applied Probability and Statistics.
- [37] M. H. ROTHKOPF AND S. S. OREN, *A closure approximation for the nonstationary $m/m/s$ queue*, Management Science, 25 (1979), pp. 522–534.
- [38] T. SIMUNIC, L. BENINI, P. GLYNN, AND G. DE MICHELI, *Event-driven power management*, IEEE TRANS. COMPUTER-AIDED DESIGN, 20 (2001), pp. 840–857.
- [39] T. SIMUNIC, H. VIKALO, P. GLYNN, AND G. DE MICHELI, *Energy efficient design of portable wireless systems*, 2000.
- [40] W. WHITT, *Stochastic-process limits*, Springer Series in Operations Research, <http://www.columbia.edu/~ww2040/supplement.html>, 2002. An introduction to stochastic-process limits and their application to queues: Internet supplement.
- [41] ———, *Stochastic-process limits*, Springer Series in Operations Research, Springer-Verlag, New York, 2002. An introduction to stochastic-process limits and their application to queues.
- [42] 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), vol. 28 of Fields Inst. Commun., Amer. Math. Soc., Providence, RI, 2000, pp. 49–71.
- [43] J. YONG AND X. Y. ZHOU, *Stochastic controls*, vol. 43 of Applications of Mathematics, Springer-Verlag, New York, 1999.