Asymptotically Optimal Control for Some Time-varying Stochastic Networks

April 9, 2008

Abstract

The present work can be placed in the realm of asymptotic stochastic optimal control in the setting of time-varying stochastic networks. More precisely, our aim is to identify controls that are asymptotically optimal (with respect to some specific, relevant performance measures) in the uniform acceleration regime (i.e., the limit of large mean arrival and service rates) and determine how well they perform for actual systems. More specifically, we consider a single station and a tandem queueing network and use the framework of strong approximations to identify asymptotically optimal controls both in the fluid regime as well as in the second-order or "diffusion" regime. The latter requires an appropriate extension to the non-stationary context of the usual notion of asymptotic optimality used for stationary networks in heavy traffic.

Contents

Chapter 1

Introduction

1.1 Introduction

1.1.1 Limiting regimes in stochastic networks

For the most part, stochastic networks are not tractable in any immediate sense (e.g., it is not possible to obtain closed-form expressions for quantities of obvious interest, for instance, the waiting time). More precisely, functions of the primitives of the system (say, measures of performance such as the sojourn time or delay) are not readily available in a closed form. Hence, an asymptotic approach towards performance analysis and optimal control of stochastic networks must be taken. Given that this thesis is aimed at the study of several control problems, we will mainly focus on the optimal control aspects of the existing literature.

The "conventional" regime In the "conventional" regime, the number of servers is kept fixed at a finite number, while space and time are scaled. One line of research has been dedicated to fluid models. Most of these results are in the context of network design for time-homogeneous systems (see [Mey01], [Mey03] and references therein). A noteworthy example of a fluid model in the time-inhomogeneous setting is [CADX04] in which the authors look at an optimal resource allocation control problem for a (stochastic) fluid model with multiple classes, where the controller dynamically schedules different classes in a system which experiences an overload over a transient period of time.

On the other hand, an immense body of work is dedicated to time-homogeneous systems in heavy traffic. Here, under appropriate assumptions on the convergence of the net arrival and service rates, the approach of the traffic intensity towards its critical level of stability 1 is ensured by speeding up time by a factor N while scaling the queue lengths down by a factor $\sqrt{N}$. Formally passing to the limit per the just described scaling, one obtains a reflected diffusion process. Then, the corresponding *Brownian control problem* (henceforth abbreviated as BCP) can be formulated, and an attempt to identify optimal or asymptotically optimal polices can be made (see, e.g., [Wil00] for references on this subject). A concise overview of the basic principles of heavy-traffic analysis and control in the time homogeneous case is given in [Har90] by the pioneer of this method M. J. Harrison.

Another asymptotic regime that has become a focus of attention in recent years is the so-called "Halfin-Whitt" regime, in which the number of servers is scaled up by an index N, while the number of customers in the queue and the number of idle servers are scaled down by either N (for fluid scaling) or $\sqrt{N}$ (for diffusion scaling). For more on control problems in this setting, one should consult [AMR04] and references therein. This thesis focuses on optimal control problems for time-inhomogeneous systems in the conventional asymptotic regime.

1.1.2 Time-dependent queueing systems

Most real-world queueing systems evolve according to laws that vary with time. A considerable body of literature has been devoted to the study of time-homogeneous models. While these models may provide reasonably good approximations for slowly varying systems, they completely fail to capture many important time-dependent phenomena such as periodicity and surges in demand. In particular, controls that are designed to optimize a specific performance measure in a stationary network may be significantly sub-optimal in the presence of significant temporal variations. It is, therefore, crucial to understand how to design optimal controls in the presence of non-stationarity. The present work takes a step in that direction.

Of course, there are exceptions to the above statements, as in the case of [Hal91]. Moreover, there is the study of the $M_t/G/\infty$ queue (i.e., the queue with arrivals modeled by a time-inhomogeneous Poisson process, i.i.d. potential service times and infinitely many servers) in [EMW93b], and its continuation in the case of sinusoidal arrival rates [EMW93a]. The latter paper also contains an application of the findings gathered in the sinusoidal case to the case of general periodic arrival rates. Furthermore, stochastic networks consisting of queues of the above type (time-varying Poisson exogenous arrival processes and infinitely many servers) have been considered in [MW93]. A succinct survey of available methods can be found in [Mas02] - a rich source of references on the analysis of time-varying queues. Although, this work is motivated by applications in telecommunications, the results are general. An alternative approach to handling the difficulties presented by time-inhomogeneity uses computer simulations. One example is [GK95] where *simple peak-hour approximation* is analysed.

Uniform acceleration The same rationale as in the time-homogeneous case discussed above brings forth consideration of asymptotic behavior as a remedy for the difficulties time-inhomogeneity presents. The type of asymptotic analysis needs to be appropriate to the model and simplify the technical aspects of the problem while retaining the features of the system that are crucial to the specific problem at hand. We first note that when studying time-inhomogeneous stochastic networks, a scaling of time, as described in the context of "conventional" heavy traffic, is not an option. In fact, such an approach would result in the homogenization of the problem and elimination of the time-dependence.

An attempt at combining the theory of queues with stationary stochastic arrivals with the deterministic theory of time-dependent queues is present in the work of Newell (see [New68a], [New68b] and [New68c]). This work was later incorporated into and expanded upon in [New71]. These papers focus on the study of a queue as it transitions through saturation (e.g., as the arrival rate increases and reaches the constant service rate) and goes back from the maximum queue length to the equilibrium (e.g., as the arrival rate decreases). In addition, a "mild rush-hour" phenomenon is explored as the same tools which work in the study of the first two transitions are not applicable here. The heuristic analysis in all these works is based on the diffusion equation (Fokker-Planck equation) for the distribution of the queue length. This strategy hints at the direction of later efforts in building a framework for asymptotics in time-varying queues. A systematic treatment similar in spirit to the intuitive one of Newell (although completely unrelated) was developed by Keller and can be found in [Kel82]. Therein, the author uses a small scaling parameter in order to derive formal asymptotic expansions of the queue length at a certain time in the $M_t/M_t/1$ setting (i.e., in the case of a single server queue with nonhomogeneous Poisson arrivals and independent exponentially distributed potential service times). Also, different phases of the system depending on its saturation are defined in this work.

Most of the work in this thesis is based on [MM95] - a rigorous treatment of the above works. There, the arrival and the potential departure rates are scaled up by N, while the queue lengths are scaled down by N. In this framework, the arrival and potential departure rates are scaled by the same factor, and the method itself is referred to as the *uniform acceleration* method. The authors of [MM95] employ the theory of strong approximations (see, e.g., [CR81] and [CH93]) to develop a Taylor-like expansion of sample paths of queue lengths, establishing a Functional Strong Law of Large Numbers and a Functional Central Limit Theorem. Furthermore, the demanding task of identifying explicit forms of the first order (in the almost sure sense) and second order (in the distributional sense) approximations of the queue lengths is accomplished. Chapter 9 of [Whi02a] relaxes some technical assumptions posited in [MM95] and exhibits some more general results. An off-shoot of the expansion of the queue length developed in [MM95] is the study of the second order approximation term in the said expansion in terms of a directional derivative of the one-sided regulation map (in an appropriate topology on the path-space). In fact, this is the basis of the exposition in [Whi02a]. This point of view inspired the development of a purely mathematical theory aimed at understanding and evaluating the directional derivatives of more general regulation maps generated by stochastic networks with more elaborate routing. The reader is directed to [MR06] for an intuitive introduction into this theory, as well as an overview of related references.

The analysis in [HHBM06] is very illustrative of the natural environment of processor sharing for the the uniform acceleration method and it introduces the important notion of sojourn time in this setting. Another point of view regarding the use of strong approximations is taken in

[Hor92]. In this paper the author studies the rate of convergence of the expected values and distributions of queue length processes when they are approximated by (reflected) Brownian motions in the context of strong approximation theory.

1.1.3 Contributions of this thesis

Our aim in the present work is to identify controls that are asymptotically optimal (with respect to some specific performance measures) in the uniform acceleration regime (i.e., the limit of large mean arrival and service rates) and determine how well they perform for actual systems. More specifically, we consider a single station and a tandem queueing network and use the framework of strong approximations to identify asymptotically optimal controls (with respect to some relevant performance measures) both in the fluid regime as well as in the second-order or "diffusion" regime. The latter requires an appropriate extension to the non-stationary context of the usual notion of asymptotic optimality used for stationary networks in heavy traffic.

We should emphasize that an important motivation behind the choice of uniform acceleration as a tool for handling the control problems presented in the sequel is the fact that this scheme keeps the ratio between the arrival rate (a parameter) and the service rate (the control) constant. This hints at the robustness of this approach with respect to the choice of scaling of actual parameters in preparation for the asymptotic analysis.

Before formulating the asymptotic optimal control problems, we highlight some unique aspects of time-inhomogeneous systems.

Flow of information Let us briefly return to optimal control in the time-homogeneous setting (say, the Brownian control problem (BCP) mentioned in Subsection 1.1.1). In that context, the only 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 that the system is in. 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.

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 on the given parameters of the model of the system. In fact, in Theorems 3.4.8, 3.4.10 and 4.8.11 we provide classes of deterministic asymptotically optimal policies for particular performance measures. However, as we can see in Section 4.7, there are specific models and sets of parameters where the optimal perfromance which can be achieved with *stochastic* (i.e., state-dependent) policies cannot be approached by a sequence of deterministic policies. Therefore, we must formulate the structure of the accumulation of information over time in dependence on the past and present states of the system. This is done in Appendix B. The formulation and solution of an asymptotic optimal control problem which requires stochastic optimal policies is carried out in Section 4.6, the main result being Theorem 4.6.23.

Performance analysis The main source of tools for performance analysis for this thesis is [MM95]. To ensure relative self-containedness, the main results of [MM95], as well as some useful consequences, are exhibited in Subsection 1.2.2.

Implementation The idea behind the implementation of the asymptotic regimes discussed above is imagining that the "actual" system is embedded in a sequence of systems approaching (in an appropriate sense) the limiting system. The question of interpretation of the results of the performance analysis and solutions to the optimal controls problems in this setting is natural. This is of particular importance in the case of optimal control, as we would like for the proposed optimal policies to be implementable in the actual system which inspired the problem in the first place. In the case of Brownian control problems (BCPs), this connection is more-or-less straightforward (see, e.g., Section 5.5 of [Whi02b] for an overview of this subject). On the other hand, in the case of time-inhomogeneous queues it is not immediately clear what the appropriate choice of the index N in the uniform acceleration corresponding to the actual system should be (see discussion in Section 2.2.1). However, note that this method preserves the ratio between the arrival and service rates. Bearing this in mind and considering Theorems 3.4.8, 3.4.10, 4.6.23 and 4.8.11, we can see that there are classes of optimal policies whose (asymptotic) performance does not depend on the choice of the index assigned to the actual system. More precisely, in the case of Theorems 3.4.8, 3.4.10 and 4.8.11 the optimal policies themselves can be constructed in a manner rendering them independent of the choice of the index. In the case of Theorem 4.6.23, it is the asymptotic performance of the proposed class that does not depend on N.

Applications The above mentioned expository paper [Mas02] considers the applications of time-varying stochastic networks to telecommunications. In the context of computer engineering, our work is closest to the fields of *power aware scheduling* and *temperature aware scheduling* (see, e.g., [BKP04]). Finally, as we will explain in some more detail when formulating the problems, the models we consider can be readily interpreted in optimization of *manufacturing systems* (see, e.g., Subsection 3.1.5).

Organization of the thesis The remainder of the present chapter contains the notation valid throughout the rest of the text and some preliminary results needed in later chapters. For the reader's convenience, all main results are gathered and displayed in Chapter 2. Also, the general philosophy behind our approach is presented and placed within the context of the existing literature. Chapter 3 is dedicated to the detailed study of the single-station control problems, while Chapter 4 deals with the tandem system. Finally, there is an Appendix containing all the auxiliary results.

Remark 1.1.1. All simulation results displayed in this thesis which involve time-inhomogeneous Poisson processes were done by discretizing (equidistantly) the rate functions. The values assigned to the discretized rates on any interval in the equidistant grid were the values of the actual rate at the left endpoint of the interval. The rates we used in the simulations were all sufficiently smooth to justify this approach.

1.2 Notation and Preliminaries

This section is dedicated to the introduction of all the notation and some preliminary results that are used repeatedly throughout the rest of the thesis.

1.2.1 Notation

Function Spaces and Mappings

- $\mathbb{L}^0_+(\Omega, \mathcal{F}, \mathbb{P})$ denotes the set of all nonnegative random variables on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$;
- $\mathbb{L}^1[0, T]$ denotes the set of all integrable functions defined on $[0, T]$;
- $\mathbb{L}^1_+[0, T]$ denotes the set of all non-negative integrable functions defined on the domain $[0, T]$;
- $\mathcal{A}$ denotes the space of all functions $F : [0, T] \to \mathbb{R}$, of the form $F(t) = \int_0^t f(s)ds$, with $f \in \mathbb{L}^1[0, T]$;
- $\mathcal{A}_+$ denotes the space of all functions $F : [0, T] \to \mathbb{R}$, of the form $F(t) = \int_0^t f(s)ds$, with $f \in \mathbb{L}^1_+[0, T]$;
- $\mathcal{I} : \mathbb{L}^1[0, T] \to \mathcal{A}$ is the (integral) mapping $\mathcal{I}(f) = \int_0^{\cdot} f(s)ds$, for $f \in \mathbb{L}^1[0, T]$;
- $\mathcal{D}$ denotes the set of all real-valued right-continuous functions on $[0, T]$ with finite left limits at all points in $(0, T]$;
- $|| \cdot ||_T$ is the uniform convergence norm on the space $\mathcal{D}$.

Processes

- N^+ and N^- are independent, unit Poisson processes defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$;
- N_1^+, N_1^- and N_2^- are independent, unit Poisson processes defined on $(\Omega, \mathcal{F}, \mathbb{P})$

1.2.2 Preliminaries

Tail Sets

For the sake of completeness, we describe the meaning of the sets $\{A_n,\, ev.\}$ (A_n happens *eventually*, i.e., for all but finitely many n) and $\{A_n,\, i.o.\}$ (A_n happens *infinitely often*), for a sequence $\{A_n\}$ of measurable sets on our probability space. These events are defined by

$$\begin{aligned} \{A_n,\, ev.\} &= \left\{ \sum_{n\geq 1} \mathbf{1}_{A_n^c} < \infty \right\} = \bigcup_{n\geq 1} \bigcap_{k\geq n} A_k, \\ \{A_n,\, i.o.\} &= \left\{ \sum_{n\geq 1} \mathbf{1}_{A_n} = \infty \right\} = \bigcap_{n\geq 1} \bigcup_{k\geq n} A_k. \end{aligned} \tag{1.2.1}$$

Reflection Maps

Definition 1.2.1. The one-dimensional (one-sided) reflection map on $\mathcal{D}$ is denoted by $\Gamma : \mathcal{D} \to \mathcal{D}$ and defined as

$$\Gamma(x) = x + l, \text{ for every } x \in \mathcal{D}, \tag{1.2.2}$$

where $l \in \mathcal{D}$ is the unique nondecreasing mapping satisfying

$$\int_0^T \Gamma(x)(t)dl(t) = 0,$$

i.e., l is "flat" on the region where $\Gamma(x) > 0$.

The mapping Γ is also referred to as the Skorokhod map on $[0, \infty)$, named after A. V. Skorokhod who introduced this mapping for continuous functions in [Sko61]. The function l has a well known explicit form

$$l(t) = \sup_{s\leq t}[-x(s)]^+, \text{ for every } t \in [0, T], \tag{1.2.3}$$

which justifies the *uniqueness* claim from the above definition (see, e.g., p. 439 of [Whi02b] or Section 2.2 of [Har90]). Moreover, for every other nondecreasing function $y \in \mathcal{D}$ satisfying $y + x \geq 0$, we have that

$$y(t) = \sup_{s\leq t}[y(s)]^+ \geq \sup_{s\leq t}[-x(s)]^+ = l(t), \text{ for every } t \in [0, T]. \tag{1.2.4}$$

In other words, the function l is the minimal element (in the partial ordering of pointwise comparison) of the set $\{y \in \mathcal{D} : x + y \geq 0\}$.

From the explicit expressions given in (1.2.2) and (1.2.3), it is easy to see that the mapping Γ is Lipschitz continuous in the uniform and M_1 topologies (see Lemma 13.5.1. and Theorem 13.5.1. in [Whi02b]).

A generalization of the Skorokhod map that constrains functions to remain in a bounded interval follows:

Definition 1.2.2. We denote the two-sided reflection map restricting functions in $\mathcal{D}$ to the region $[0, K]$ by $\Gamma^K : \mathcal{D} \to \mathcal{D}$ and define it by

$$\Gamma^K(x) = q = x + l - u, \text{ for all } x \in \mathcal{D}, \tag{1.2.5}$$

where (q, l, u) is the *unique* triplet in $\mathcal{D}^3$ such that

(i) the functions $l, u \in \mathcal{D}$ are nondecreasing;

(ii) l is "flat" on the region where $q > 0$, i.e., $\int_0^T q(t)\, dl(t) = 0$;

(iii) u is "flat" on the region where $q < K$, i.e., $\int_0^T (q(t) - K)\, du(t) = 0$.

We will refer to l and u as the regulator functions associated with x and K.

In the terminology of Definition 1.2 of [KLRS06] and following our present notation, one would say that the pair $(q, l - u)$ solves the Skorokhod problem on $[0, K]$ for x (see also Section 14.8 of [Whi02b] for a discussion of the two-sided reflection map).

The lower regulator map l and the upper regulator map u of Definition 1.2.2 can be tied together through the following well-known pair of identities

$$l(t) = \sup_{s \le t}[-x(s) + u(s)]^+,\ u(t) = \sup_{s \le t}[x(s) - l(s) - K]^+, \text{ for every } t. \tag{1.2.6}$$

The above identities are mentioned as equation (1.9) in [KLRS06], and are justified in Section 14.8 of [Whi02b] and Section 2.4 of [Har90].

Next, we present the useful minimality feature of the regulators l and u from (1.2.5), analogous to the one expressed earlier (see (1.2.4)) in the context of the one-sided reflection map.

Proposition 1.2.3. *Given $K \in (0, \infty)$ and $x \in \mathcal{D}$ such that $x(0) = 0$, let (q, l, u) be the unique triplet associated with K and x as specified in Definition 1.2.2. Then, for all nonnegative nondecreasing l' and u' such that the function $q' = x + l' - u'$ is constrained within $[0, K]$, we have that $l \le l'$ and $u \le u'$.*

Proof. Let (q, l, u) and (q', l', u') be as in the statement of the lemma. We define the following sequence[1] of instances in $[0, T]$

$$\begin{aligned} t_0 &= 0, \\ t_1 &= \inf\{t \geq 0 : q(t) = K\} \wedge T, \\ t_{2i} &= \inf\{t \geq t_{2i-1} : q(t) = 0\} \wedge T, \text{ for every } i \in \mathbb{N}, \\ t_{2i+1} &= \inf\{t \geq t_{2i} : q(t) = K\} \wedge T, \text{ for every } i \in \mathbb{N}. \end{aligned}$$

Claim: There exists an index J such that $t_j = T$, for every $j \geq J$. Let us suppose that, contrary to the posited claim, there exist countably many excursions of the path q across the strip $[0, K]$. Since the sequence $\{t_j\}$ is nondecreasing and bounded from above by T, it must be convergent. Let us denote its limit by ξ.

Since $q \in \mathcal{D}$, it has left limits at all points in $[0, T]$. In particular, let us denote by ζ the left limit of q at ξ, i.e., let $\zeta = q(\xi-)$. By definition, then, for an arbitrary positive constant ε, there exists another positive constant δ such that for every $t > \xi - \delta$, we have that $|q(t) - \zeta| < \varepsilon$. On the other hand, since $t_j \to \xi$, there exists an index j_δ such that for every $j > j_\delta$ we have that $t_j > \xi - \delta$. Thus, for every $j > j_\delta$, it must be that $|q(t_j) - \zeta| < \varepsilon$. Choosing $\varepsilon = \frac{K}{5}$ produces a contradiction with the definition of the sequence $\{t_j\}$.

"Mathematical Induction" We continue with the proof of the main claim. The strategy is to prove the claim in an "inductive" manner starting with time 0 and progressing towards time T considering one interval of the form $[t_{j-1}, t_j)$, $1 \leq j \leq J$, at a time.

By definition, we have that $u(t) = 0$ for every $t \in [0, t_1)$. Hence, we can conclude that $u \leq u'$ on $[0, t_1)$. At the same time, due to the first identity in (1.2.6), we deduce that

$$l(t) = \sup_{s \leq t}[-x(s) + u(s)]^+ = \sup_{s \leq t}[-x(s)]^+.$$

As q' and u' are both nonnegative, we have that $x' + l' \geq 0$. Reiterating the argument of (1.2.4), we get that $l \leq l'$ on $[0, t_1)$.

Next, we focus on the interval $[t_1, t_2)$. Since the process q does not visit 0 during this interval, we immediately get that the function l remains constant, i.e., $l(t) = l(t_1)$, for every $t \in [t_1, t_2)$. As we have already proven that $l(t_1) \leq l'(t_1)$ and since the function l' is nondecreasing, we deduce that $l(t) \leq l'(t)$, for every $t \in [t_1, t_2)$.

Simultaneously, by the second equality in (1.2.6), we have that

$$u(t) = \sup_{s \leq t}[x(s) + l(s) - K]^+ \leq \sup_{s \leq t}[q'(s) + u'(s) - K]^+, \text{ for every } t \in [t_1, t_2).$$

Since $q' \leq K$ and u' is, by assumption, nondecreasing, the last display gives us that $u \leq u'$ on $[t_1, t_2)$.

[1]Since $x \in \mathcal{D}$, the possibility of uncountably many instances of the type described below is immediately ruled out.

So far we proved that the announced minimality of l and u holds true on $[0, t_2)$. In the spirit of the principle of mathematical induction, let us assume that the minimality property holds on $[0, t_{2k})$, for some $k \geq 1$. We aim to prove that the analogous claim holds on $[t_{2k}, t_{2k+2})$.

On $[t_{2k}, t_{2k+1})$, we have that $q < K$, so there is no need for upper regulation during this period, i.e., $u(t) = u(t_{2k})$, for every $t \in [t_{2k}, t_{2k+1})$. Consequently, using once more the first identity in (1.2.6) and the inductive hypothesis, we arrive at

$$l(t) = \sup_{s \leq t}[-x(s) + u(s)]^+ \leq \sup_{s \leq t}[-x(s) + u'(s)]^+ = \sup_{s \leq t}[-q'(s) + l'(s)]^+$$

for every $t \in [t_{2k}, t_{2k+1}]$. Since q' is, by construction, nonnegative and the function l' is nondecreasing, we get $l \leq l'$ on the entire interval $[0, t_{2k+1})$.

Finally, we focus on the interval $[t_{2k-1}, t_{2k})$. On this interval, the function q is always strictly positive, which implies that $l(t) = l(t_{2k-1})$, for every $t \in [t_{2k-1}, t_{2k})$. As above, this straightforwardly implies that $l \leq l'$ on $[t_{2k-1}, t_{2k})$. On the other hand, thanks to the second equality in (1.2.6) and the facts that $q' \leq K$ and u' in nondecreasing, the upper regulator u satisfies the following chain of (in)equalities

$$u(t) = \sup_{s \leq t}[x(s) + l(s) - K]^+ \leq \sup_{s \leq t}[q'(s) + u'(s) - K]^+ \leq \sup_{s \leq t}[u'(s)]^+ = u'(t),$$

for every $t \in [t_{2k+1}, t_{2k})$. □

The following consequence of Theorem 4.2 of [BW92] will be used extensively in the sequel. We simply restate the result using the notation we introduced above.

Proposition 1.2.4. *Let $\{x_n\}$ be a sequence in $\mathcal{D}$ converging uniformly to a function $x \in \mathcal{D}$ and let the sequence $\{(q_n, l_n, u_n)\}$ denote the triplets uniquely determined by the two-sided regulator map applied to the terms in the sequence $\{x_n\}$. Then we have that $(q_n, l_n, u_n) \to (q, l, u)$, in the uniform topology, where (q, l, u) is the triplet generated by the two-sided regulator map applied to the function x.*

Limit Theorems

The next result is a by-product of the proof of Theorem 2.1 from [MM95].[2]

Theorem 1.2.5 (FSLLN)**.** *Let λ and μ be nonnegative, integrable functions on $[0, T]$ and let N^+ and N^- be independent unit Poisson processes on a common probability space. Define the sequence of stochastic processes $\{X^{(n)}\}$ as*

$$X^{(n)} = N^+(n\mathcal{I}(\lambda)) - N^-(n\mathcal{I}(\mu)). \tag{1.2.7}$$

[2]In the cited paper, the authors have a standing assumption that μ is strictly positive - an assumption we are not making anywhere. However, the proof of this particular claim presented in their paper does not use the strict positivity assumption.

Figure 1.1: Runs of a single station with rates $\lambda(t) = 1 + \cos(2\pi t)$ and $\mu = \frac{1}{2}\lambda$. The uniform acceleration indices are $n = 50, 100, 500$.

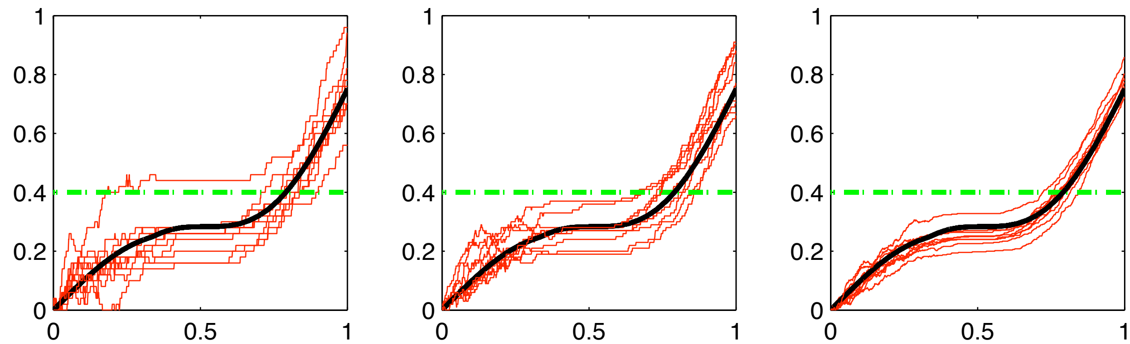

Then we have that in the uniform topology

$$\frac{1}{n}X^{(n)} \to \mathcal{I}(\lambda - \mu), a.s. \tag{1.2.8}$$

In the rest of the text, we will refer to the last theorem as the Functional Strong Law of Large Numbers or FSLLN. An illustration of Theorem 1.2.5 by way of simulations is given in Figure 1.1.

The following corollary is a simple consequence of Theorem 1.2.5 and Proposition 1.2.4.

Corollary 1.2.6. *Suppose that the sequence of random processes $\{X^{(n)}\}$ is defined as in (1.2.7). Let the sequence of triplets of stochastic processes obtained through the application of the two-sided reflection map to the terms in the said sequence be denoted by $\{(Q^{(n)}, L^{(n)}, U^{(n)})\}$ and let (q, l, u) be the triplet obtained by applying the two-sided reflection mapping of Definition 1.2.2 to the function $x = \mathcal{I}(\lambda - \mu)$. Then we have that as $n \to \infty$,*

$$\left(\frac{1}{n}Q^{(n)}, \frac{1}{n}L^{(n)}, \frac{1}{n}U^{(n)}\right) \to (q, l, u), \ a.s., \tag{1.2.9}$$

in the uniform topology.

The following result addresses the second order approximations of the queue length processes. It will be referred to as the Functional Central Limit Theorem or FCLT.

Theorem 1.2.7 (FCLT)**.** *Suppose that the functions λ and μ, as well as processes $\{X^{(n)}\}$ are as in Theorem 1.2.5. Let the sequence of stochastic processes $\{Q^{(n)}\}$ be given as $Q^{(n)} = \Gamma(X^{(n)})$, and let $\bar{q} = \Gamma(\mathcal{I}(\lambda - \mu))$. Then we have that*

$$\sqrt{n}\left(\frac{1}{n}Q^{(n)} - \bar{q}\right) \Rightarrow \hat{Q},$$

in distribution, with respect to the M_1 topology, where for every $t \in [0, T]$ the process $\hat{Q}$ is determined by

$$\hat{Q}_t \stackrel{(d)}{=} W(\mathcal{I}_t(\lambda + \mu)) + \sup_{s \in \Phi_{-\mathcal{I}(\lambda-\mu)}(t)} [-W(\mathcal{I}_s(\lambda + \mu))],$$

with

$$\Phi_\xi(t) = \{s \leq t : \sup_{u \leq t}[\xi(u)] = \xi(s)\},$$

for any $\xi \in \mathcal{C}[0, T]$ and where W is a standard Brownian motion.

The origin of the above result is [MM95]. There, it was stated under the additional restriction that the process $\hat{Q}$ have only a finite number of discontinuities on any compact set. However, this condition can be done away with as is shown, e.g., in [Whi02a] and [MR06]. We illustrate the FCLT in Figure 1.2. The first box displays the fluid limits of the arrival and departure rates, the netput and the queue length. The vertical lines represent the partition of $[0, T]$ into the regions where the system is:

- overloaded, i.e., where the fluid limit of the queue length strictly positive;
- critical, i.e., where the fluid limit of the queue length is at 0 and there is no upward pushing in the Skorokhod map;
- underloaded, i.e., where the fluid limit of the queue length is at 0 and there is upward pushing in the Skorokhod map.

The second box displays the results of several simulations of the scaled queue length and the second order approximation processes

$$\bar{q} + \frac{1}{\sqrt{n}}\hat{Q}$$

with the index of acceleration $n = 50$. Finally, the third box shows the process $\hat{Q}$ itself.

Both the FSLLN and the FCLT are consequences of the Strong Approximation Theorem which we state next in the form in which it was exhibited in [MM95].

Theorem 1.2.8 (Strong Approximation)**.** *The sequence of stochastic processes $\{Q^{(n)}\}$ can be realized on a common probability space, also supporting two independent standard Brownian motions W^+ and W^- in a way that*

$$Q^{(n)} = \Gamma(\tilde{X}^{(n)}) + O(\log(n)),$$

where

$$\tilde{X}^{(n)} = n\mathcal{I}(\lambda - \mu) + W^+(n\mathcal{I}(\lambda)) - W^-(n\mathcal{I}(\mu)).$$

Figure 1.2: Functional Central Limit Theorem

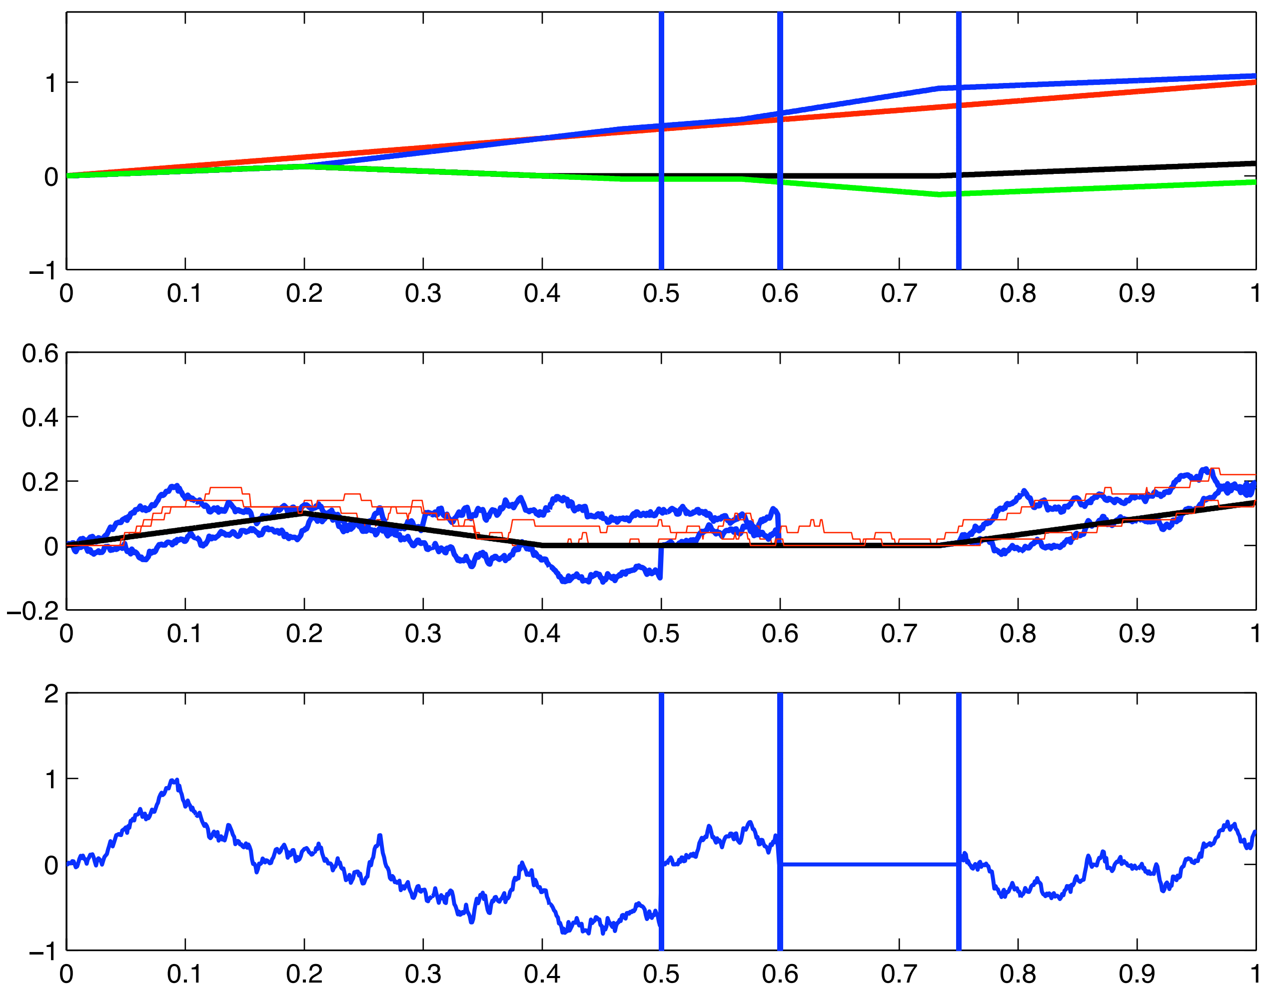

Chapter 2

Main Results

2.1 Outline

This thesis is dedicated to a detailed study of a family of optimal control problems of time-inhomogeneous stochastic networks. The emphasis in this work is on the methodology, so the network aspects (network topologies and stochastic models), as well as the performance aspects (forms of the performance measures and controls) are kept simple. In the present section, we briefly describe all optimal control problems which appear in this thesis.

As a mnemonic aid for the reader going through the descriptions of various optimal control problems, we provide a tableau which encompasses all the main results (see Figure 2.1). Each cell in the table, except for the shaded right cell in the second row, corresponds to an optimal control problem that is considered in the thesis. The reader is advised to consult Subsection 1.2.1 for notation and standing assumptions for the entire thesis.

2.1.1 The General Philosophy

In the next chapter we describe several control problems associated with time-inhomogeneous queueing networks. In each case, an exact analytic solution is not feasible. Thus, we embed the actual system into a sequence of systems with arrival rates (and buffer capacities, where applicable) tending to infinity, and identify a sequence of controls which are asymptotically optimal in the sense described precisely below. This sequence of systems is constructed by means of the so-called "uniform acceleration" scaling appropriate as the "pre-limit sequence".

In many cases, the identification of the class of asymptotically optimal sequence 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 taking FSLLN limits of the original systems, and thus usually lead to the formulation of deterministic control problems. Second-order problems, on the other hand, also take into account certain fluctuations around the FSSLN limits. In the case of a time-homogeneous queueing system, the second-

Figure 2.1: The Layout of Control Problems

		Performance Measure	
		Time above threshold (Infinite buffers)	Lost jobs (Finite buffers)
Single Station		Fluid Limit	Fluid Limit
		Second Order Approximation	
		Pre-limit	Pre-limit
The Tandem Network		Fluid Limit	Fluid Limit
		Pre-limit	Pre-limit

order approximation of a queueing network is usually given by a reflected diffusion. So, in this case, the second-order approximation of the original control problem leads to a single reflected diffusion control problem. The methodology of using fluid and diffusion control problems to identify asymptotically optimal controls is fairly well-developed in the time-homogeneous setting (see, for example, [DR00, Mey01, Mey03, Dup03] for the use of fluid control problems and [Har90, HVM97, AHS05, Kus01] for the use of diffusion control problems). In contrast, in the time-inhomogeneous case, there is relatively little rigorous work in this domain (see, e.g., [New71] for some heuristics). Indeed, one of the main aims of this thesis is to take a step towards developing a suitable methodology for optimal control in the time-inhomogeneous case. As shown in Theorem 1.2.7, in the time-inhomogeneous case, the second-order approximation of the pre-limit sequence is given by another sequence of simpler processes, rather than a single process. Thus, in this case, the second-order approximation of the pre-limit sequence of control problems leads to another sequence of control problems for which an asymptotically optimal sequence of controls needs to be determined. The second-order approximation is useful only when the second-order sequence of control problems proves easier to analyze than the pre-limit sequence of control problems. In the following sections, we describe the fluid limit, second-order (when applicable) and pre-limit versions of several control problems related to time-inhomogeneous queueing networks. In Section 2.2 we consider a single station and examine two different performance measures in Subsections 2.2.1 and 2.2.2, while in Section 2.3 we consider a tandem network and again consider two different performance measures in Subsections 2.3.1 and 2.3.2.

Before delving into the descriptions of particular optimal control problems, let us pause to ask some general questions regarding the rationale discussed above.

(i) How does one formulate the second-order control problem (the control problem analogous to the diffusion control problem in the time-homogeneous case) for time-inhomogeneous systems?

(ii) Does the consideration of the second-order approximations and the corresponding optimal control problem(s) shed more light on the optimal control of the pre-limit systems than the first-order approximation does alone?

(iii) We emphasize again that the concept of *deterministic* control that is based on the mean time-varying arrival rates (possibly gathered through observations of the system in the past), but is not dependent on the actual states of the system in the actual (controlled) run of the system, arises exclusively in the time-inhomogeneous setting. Is it possible to provide deterministic optimal controls, or does one need to resort to stochastic controls that are dynamically adjusted depending on the state of the system?

(iv) How does one apply the results in the context of asymptotically optimal control problems to the actual system? More precisely, as we will see in more detail below (see Subsection 2.2.1), the embedding of the actual system into a sequence of uniformly accelerated systems demands a choice of the "constant of acceleration" of the actual system. Thus, the above question may be understood as the following one: In which way does the choice of the "constant of acceleration" affect the choice of the control one proposes for the actual system and/or the performance of the chosen control in a run of the actual system?

The above issues are discussed in Section 2.4 in the context of specific optimal control problems considered in this thesis.

2.2 Single Station

2.2.1 The Finite Buffer Context

We consider a single station that has a finite buffer capacity K^A and has one server serving one job class on a finite time-horizon $[0, T]$. The exogenous arrival rate is modeled as an inhomogeneous Poisson process with rate λ^A, where λ^A (the mean exogenous arrival rate) is a deterministic, non-negative and integrable function[1]. The (potential) service process is assumed

[1] The superscript A is used in this section to emphasize that the parameter λ^A is associated with an actual system. As discussed below, in our analysis, the actual system will be embedded in a sequence of systems with

Figure 2.2: Single Station with a Finite Buffer

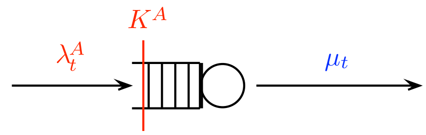

to be another Poisson process that is independent of the arrival Poisson process and has rate μ^A, where μ^A is allowed to be random, assuming values in $\mathbb{L}^1_+[0,T]$, but is non-anticipating (see Chapter B for a precise definition) and satisfies $\mathcal{I}(\mu^A)_T \leq m^A$. Note that when μ^A is deterministic, the latter inequality translates into a bound on the mean total (potential) service provided in the interval $[0,T]$. Any such service rate μ^A will be referred to as an admissible control. The system described above is depicted in Figure 2.2.

The finiteness of the buffer may cause a loss of jobs from the system. The control problem of interest is to find an optimal admissible (potential) service rate μ^A that minimizes the number of jobs lost from the system during the interval $[0,T]$. Since an exact analytic solution of this problem is not feasible, we embed the actual system into a sequence of systems with arrival rates and buffer capacities tending to infinity, and identify a sequence of controls that are asymptotically optimal (in the sense stated precisely in Definition 3.4.2). As described below, this sequence of systems is obtained by uniformly accelerating the arrival rates and service rates (see the discussion on the uniform acceleration scaling given in Subsection 1.2.2).

For any constant $N \in \mathbb{N}$, let λ, K and m be such that $K^A = NK$, $\lambda^A = N\lambda$ and $m^A = Nm$. The sequence of systems is then defined as follows. For each $n \in \mathbb{N}$ and an admissible service discipline $\mu \in \mathbb{L}^1_+[0,T]$ (which now satisfies $\mathcal{I}_T(\mu) \leq nm$), consider a station with buffer capacity nK, the exogenous arrival process given by $N^+(n\mathcal{I}(\lambda))$ and the potential service process given by $N^-(n\mathcal{I}(\mu))$, where N^+ and N^- are independent Poisson processes. In the nth system, the length of the queue then equals

$$Q^{(n)}(\mu, K) = \Gamma^{nK}(X^{(n)}(\mu)),$$

where Γ^{nK} is the two-sided reflection map introduced in Definition 1.2.2 and $X^{(n)}$ is the so-called netput process defined by

$$X^{(n)}(\mu) = N^+(n\mathcal{I}(\lambda)) - N^-(n\mathcal{I}(\mu)). \qquad (2.2.1)$$

Our goal is then to identify the sequence $\{\mu_n\}$ of admissible controls that (in the asymptotic limit as $n \to \infty$) minimizes the number of jobs lost in the n^{th} system due to finiteness of the

buffer capacities and arrival rates tending to infinity. The superscript is used here to show how the asymptotic results can be used to make inferences about an actual system with a given finite buffer capacity K^A and arrival rate λ^A. However, the superscript will be omitted in subsequent sections for ease of notation.

buffer. This is described below as the Pre-limit Sequence problem. The identification of the optimal sequence for the pre-limit problem is facilitated by first solving a simpler, deterministic control problem related to the so-called fluid or FSSLN limit of this sequence (described in Theorem 1.2.5). This is stated as the fluid limit problem below.

The Fluid Limit [***Fully covered in Subsection 3.2.2.***]

The Model We look at the process

$$\bar{Q}(\mu, K) = \bar{X}(\mu) + \bar{L}(\mu, K) - \bar{U}(\mu, K), \tag{2.2.2}$$

with $\bar{X}(\mu) = \mathcal{I}(\lambda - \mu)$ and where $\bar{L}(\mu, K)$ and $\bar{U}(\mu, K)$, as in Definition 1.2.2, are the regulators associated with $\bar{X}(\mu)$ and K.

The Performance Measure The fluid-limit version of the performance measure corresponding to the jobs lost due to the finiteness of the buffer is

$$\bar{U}_T(\mu, K).$$

The Result The entire class of policies that are optimal for the above control problem is fully described in Subsection 3.2.2. A short description of fluid-optimal service disciplines is that they do **not** cause

- upward pushing during the entire interval $[0, T]$, i.e., $\bar{L}_T(\mu, K) = 0$ for all fluid-optimal μ;
- downward pushing until time $\tau(K + m) = \inf\{t \in [0, T] : \mathcal{I}_t(\lambda) > K_1 + m\}$.

In the case that $\mathcal{I}_T(\lambda) > K + m$, this is an equivalence result, otherwise the above conditions are merely sufficient for fluid-optimality. One representative fluid-optimal policy is $\mu^* = \lambda \mathbf{1}_{[0,\tau(m)]}$.

The Pre-limit Sequence [***Fully covered in Subsection 3.4.2.***]

The model Based on Definition 1.2.2, we can express the queue length in the n^{th} system as

$$Q^{(n)}(\mu, K) = X^{(n)}(\mu) + L^{(n)}(\mu, K) - U^{(n)}(\mu, K),$$

where random processes $L^{(n)}(\mu, K)$ and $U^{(n)}(\mu, K)$ are the regulators associated with $X^{(n)}(\mu)$ and nK as per Definition 1.2.2.

The Performance Measure The number of jobs lost from the system over the time interval $[0, T]$ is recorded in the amount of regulation needed to restrict the netput process below the level K, i.e., the number of lost jobs when the service discipline μ is used is exactly the random variable $U_T^{(n)}(\mu, K)$. Upon necessary normalization, the quantity we wish to minimize (in the limit, almost surely) across controls μ becomes $\frac{1}{n} U_T^{(n)}(\mu, K)$.

Figure 2.3: Sequence of Single Stations with Infinite Buffers

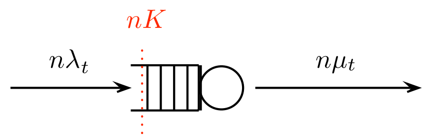

The Result Here we succeed in obtaining a class of deterministic control sequences which are asymptotically optimal in the sense of Definition 3.4.2. In an informal way, this class can be described as containing all sequences of admissible deterministic service disciplines $\{\mu_n\}$ such that the "mean" of the queue length, given by $\Gamma^K(\mathcal{I}(\lambda - \mu_n))$, remains "a bit more than" $o(\frac{1}{\sqrt{n}})$ away from the boundaries of the strip $[0, K]$ as long as the constraint on the total amount of service allows it to. This simple rule enables the controller to avoid both upward and downward pushing for all but finitely many n in the almost sure sense, until the amount of service rendered reaches the imposed constraint.

2.2.2 The Infinite Buffer Context

Having discussed the rationale for doing so in the context of the finite buffer, we immediately proceed to the sequence of systems constructed through the uniform acceleration procedure. In this case, the queue lengths are given by

$$Q^{(n)}(\mu) = \Gamma(X^{(n)}(\mu)), \tag{2.2.3}$$

for every admissible μ (again, see Chapter B for the definition of admissibility), and where $X^{(n)}(\mu)$ is the netput process from (2.2.1) and Γ the one-sided reflection map of Definition 1.2.1. A schematic description of one system in this sequence can be seen in Figure 2.3.

The positive constants nK now represent thresholds and the performance of a control is measured in terms of the time the buffer exceeds that level. In other words, the penalty incurred in a run of the system equals the amount of time the queue length spends above the level nK.

The Fluid Limit [***Fully covered in Subsection 3.2.1.***]

The Model The fluid-limit process in this case is

$$\bar{Q}(\mu) = \bar{X}(\mu) + \bar{L}(\mu),$$

where $\bar{X}(\mu) = \mathcal{I}(\lambda - \mu)$ and $\bar{L}(\mu)$ emerges from decomposition (1.2.2) in Definition 1.2.1.

Figure 2.4: Fluid arrival rates, departure rates and limits of the queue lengths for three fluid-optimal policies, where $\lambda(t) = 1 + \cos(2\pi t)$, $K = 0.4$ and $m = 0.25$.

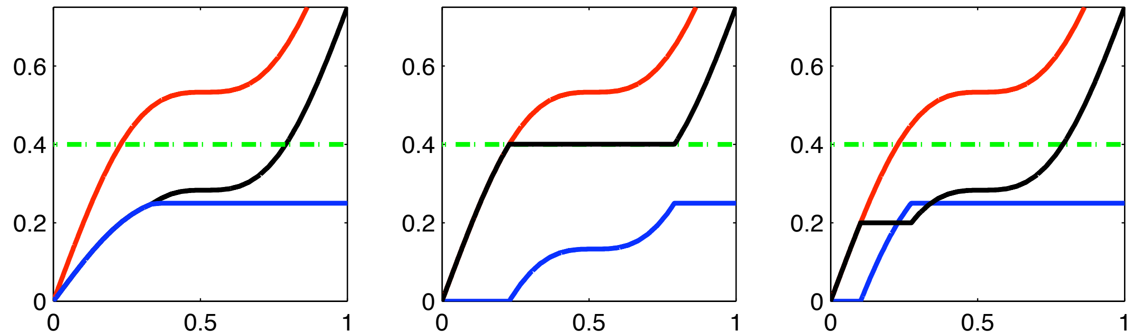

The Performance Measure An admissible policy μ causes the penalty of

$$\int_0^T \mathbf{1}_{\{\bar{Q}_t(\mu)>K\}}\, dt.$$

The Result The class of fluid-optimal controls is identical to the one in the fluid-limit analysis in the finite buffer case above.

Three different fluid-optimal policies are displayed in Figure 2.4. In the three graphs, the function which dominates the other two is the arrival rate, the functions which is (most of the time) dominated by the other two is the departure rate (i.e., a fluid-optimal policy), and the remaining curve represents the fluid limit of the queue length with the stated rates. In the order in which the fluid limits of the queue lengths appear in the figure, the policies are defined as follows:

(i) $\mu = \frac{1}{2}\lambda$ until $\mathcal{I}(\mu)$ reaches the constraint m;

(ii) $\mu = 0$ until $\mathcal{I}(\lambda)$ reaches the threshold; then, $\mu = \lambda$ until $\mathcal{I}(\mu)$ reaches the constraint m;

(iii) $\mu = 0$ until $\mathcal{I}(\lambda)$ covers half the distance to the threshold; then, $\mu = \lambda$ until $\mathcal{I}(\mu)$ reaches the constraint m.

Second Order Approximations [*Fully covered in Section 3.3*]

We now return our attention to the sequence of processes $\{Q^{(n)}(\mu)\}$ introduced in (2.2.3), for any $\mu \in \mathbb{L}^1_+[0,T]$. Theorems 2.1 and 2.2 in [MM95] or Theorem 9.6.2 in [Whi02a] both deliver the asymptotic expansion of the form

$$\frac{1}{n}Q_t^{(n)}(\mu) \overset{(d)}{=} \bar{Q}_t(\mu) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu) + o\left(\frac{1}{\sqrt{n}}\right). \tag{2.2.4}$$

In fact, this expression is given as equation (2.7) in [MM95]. The processes

$$\bar{Q}(\mu) + \frac{1}{\sqrt{n}}\hat{Q}(\mu) \tag{2.2.5}$$

can be interpreted as *second order approximations* of the processes $\frac{1}{n}Q^{(n)}(\mu)$ in a *distributional* sense.

In view of (2.2.4), we can relate the amount of time the queue length in the n^{th} pre-limit system modeled by (2.2.3) spends above the threshold K to the amount of time the random process on the right-hand side of (2.2.4) spends above K. More precisely, we consider a sequence of performance measures of the form

$$\text{meas}\left\{t \in [0,T] : \bar{Q}_t(\mu) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu) > K\right\}. \tag{2.2.6}$$

It is reasonable to expect that asymptotically optimal results on second order optimality in the sense of minimization of the expected value of the expression in (2.2.6) across deterministic service disciplines will provide insight in asymptotic performance of the pre-limit sequence.

The Model For any deterministic control μ we introduce the random process

$$\hat{Q}(\mu) = W(\mathcal{I}(\lambda+\mu)) + \sup_{s\in\Phi_{-\bar{X}(\mu)}(\cdot)}[-W(\mathcal{I}_s(\lambda+\mu))],$$

where W is a standard Brownian motion and

$$\Phi_{-\bar{X}(\mu)}(t) = \left\{s \le t : -\bar{X}_s(\mu) = \sup_{u\le t}[-\bar{X}_u(\mu)]\right\}, \text{ for all } t \in [0,T],$$

with $\bar{X}(\mu) = \mathcal{I}(\lambda-\mu)$.

The Performance Measure The controller wishes to minimize the expected value of

$$\text{meas}\left\{t \in [0,T] : \bar{Q}_t(\mu) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu) > K\right\}$$

in an asymptotic sense and across all admissible deterministic controls μ.

The Result The second order optimal class we identify contains all deterministic service disciplines $\{\mu_n\}$ such that the "mean" of the queue length $\Gamma(\mathcal{I}(\lambda-\mu_n))$ is "more than" $o(\frac{1}{\sqrt{n}})$ away from the boundaries of the strip $[0,K]$ (until the allowed amount of service is used up).[2]

[2] The exact nature of the difference between the two classes, which we attempt to express by means of the imprecise wordings "more than" and "a bit more than" modifying $o(\frac{1}{\sqrt{n}})$, will be exhibited in the main text.

Figure 2.5: Fluid limit of the queue length and several runs of the second order approximation process of (2.2.5) for $n = 100$ with $\lambda(t) = 1 + \cos(2\pi t)$ and $\mu(t) = \lambda(t)$, for t such that $\mathcal{I}_t(\lambda) \in [\frac{K}{2}, \frac{K}{2} + m]$, where $K = 0.4$ and $m = 0.25$.

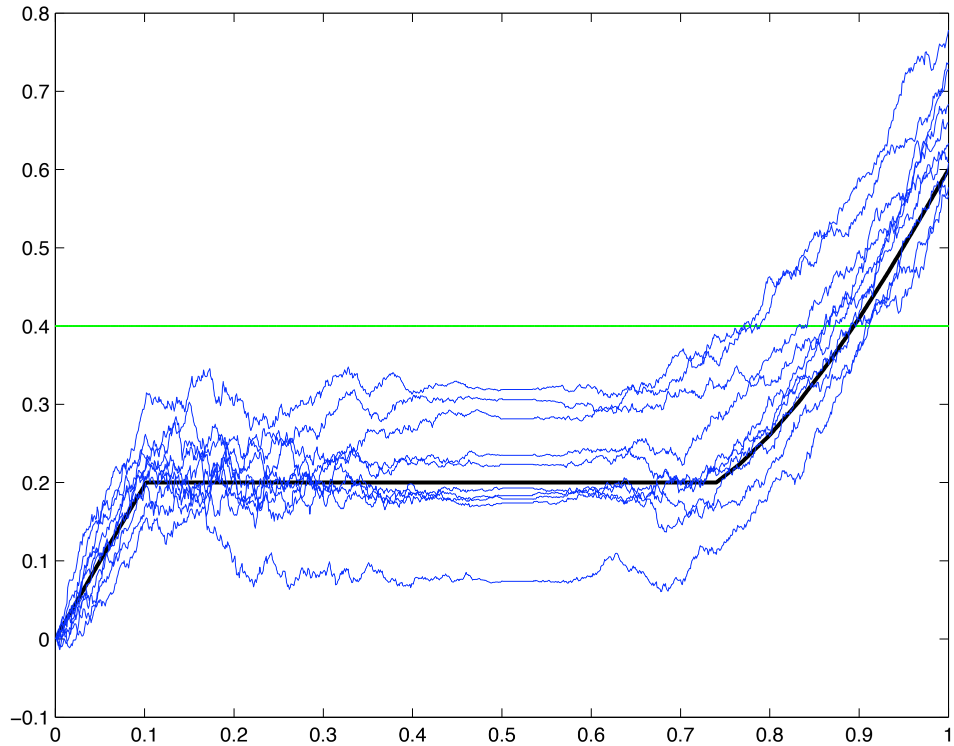

Also, the class of second order optimal sequences is a superclass of the class of all asymptotically optimal sequences for the control problem involving the pre-limit sequence of systems with finite buffers. In other words, all sequences of policies that we identified as asymptotically optimal for the pre-limit sequence of systems with finite buffers will also be second order optimal in the infinite-buffer context. Figure 2.5 illustrates this result.

The Pre-limit Sequence [*Fully covered in Subsection 3.4.1*]

The model The queue lengths are given by $Q^{(n)}(\mu) = \Gamma(X^{(n)}(\mu))$.

The Performance Measure The controller wishes to (asymptotically) minimize, in the almost sure sense, the value

$$\int_0^T \mathbf{1}_{\{\frac{1}{n}Q_t^{(n)}(\mu) > K\}}\, dt$$

by varying the admissible control μ.

Figure 2.6: Fluid limit of the queue length and several runs of the scaled pre-limit queue lengths with $\lambda(t) = 1 + \cos(2\pi t)$ and $\mu(t) = \lambda(t)$, for t such that $\mathcal{I}_t(\lambda) \in [\frac{3K}{4}, \frac{3K}{4} + m]$, where $K = 0.4$ and $m = 0.25$.

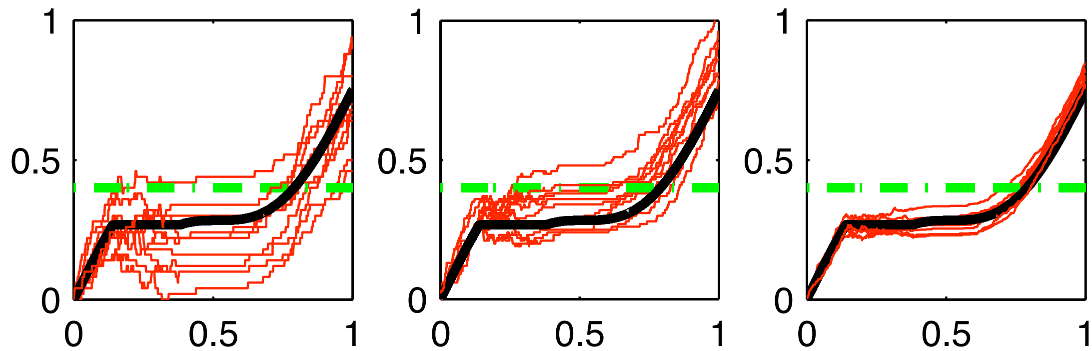

The Result The class of asymptotically optimal policies we provide is the same as the class described in the finite buffer case. Several simulated runs of the queue length for an asymptotically optimal policy are displayed in Figure 2.6.

2.3 The Tandem Network

2.3.1 The Finite Buffer Context

We consider a two-station tandem queueing network, with each station having a single server and a finite buffer capacity. The process of exogenous arrivals into the first station and the processes of potential services from both stations are independent, nonhomogeneous Poisson processes. The capacities of the two finite buffers are denoted by $K_1, K_2 \in \mathbb{R}_+ \cup \{\infty\}$. The system is shown in Figure 2.7. The controller's aim is to minimize the number of jobs lost in both stations due to the finiteness of the two buffers by means of varying the service rate at the first station among all non-anticipating services satisfying a given constraint on the total amount of service available.

Figure 2.7: The Tandem System - Finite Buffers

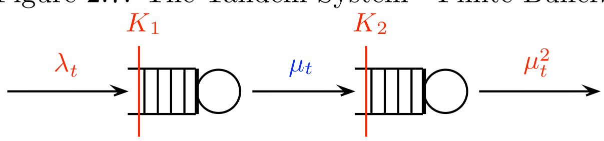

The Fluid Limit [***Fully covered in Section 4.5***]

The Model The fluid-limit versions of the queue lengths in the system depicted in Figure 2.7 and described above are

$$\bar{Q}^{(1)}(\mu, K_1) = \mathcal{I}(\lambda - \mu) + \bar{L}^{(1)}(\mu, K_1) - \bar{U}^{(1)}(\mu, K_1),$$
$$\bar{Q}^{(2)}(\mu, K_2) = \mathcal{I}(\mu - \mu^2) - \bar{L}^{(1)}(\mu, K_1) + \bar{L}^{(2)}(\mu, K_2) - \bar{U}^{(2)}(\mu, K_2),$$

where $\bar{L}^{(1)}(\mu, K_1)$ and $\bar{U}^{(1)}(\mu, K_1)$ are the regulators associated with $\mathcal{I}(\lambda - \mu)$ and K_1 through Definition 1.2.2, and $\bar{L}^{(2)}(\mu, K_2)$ and $\bar{U}^{(2)}(\mu, K_2)$ are the regulators associated with the function $\mathcal{I}(\mu - \mu^2) - \bar{L}^{(1)}(\mu, K_1)$ and K_2 per Definition 1.2.2.

The Performance Measure The goal is to minimize

$$\bar{U}_T^{(1)}(\mu, K_1) + \bar{U}_T^{(2)}(\mu, K_2)$$

over all admissible deterministic μ.

The Result We show that the policy which promptly serves all customers that arrive into the first station, as long as the resulting departures from the first station do not cause downward pushing in the second station and until the given upper bound on the amount of service is met is a fluid-optimal policy. We shall refer to this specific policy as $\mu^{*,F}$.

The Pre-limit Sequence [***Fully covered in Section 4.8***]

The Model The uniform acceleration procedure in this case gives rise to the following sequence of pairs of queue lengths

$$Q^{(1,n)}(\mu, K_1) = N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu)) + L^{(1,n)}(\mu, K_1) - U^{(1,n)}(\mu, K_1),$$

with $L^{(1,n)}(\mu, K_1)$ and $U^{(1,n)}(\mu, K_1)$ representing the regulator maps associated with the process $N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu))$ and the constant nK_1 per Definition 1.2.2, and

$$Q^{(2,n)}(\mu, K_2) = N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu, K_1) - N_2^-(n\mathcal{I}(\mu^2)) + L^{(2,n)}(\mu, K_2) - U^{(2,n)}(\mu, K_2),$$

where the regulator maps $L^{(2,n)}(\mu, K_2)$ and $U^{(2,n)}(\mu, K_2)$ are associated with $N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu, K_1) - N_2^-(n\mathcal{I}(\mu^2))$ and nK_2 according to Definition 1.2.2.

Figure 2.8: The Tandem system - Infinite buffers

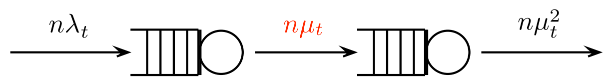

The Performance Measure For every index n, the penalty associated with the admissible service discipline μ is the expected value of

$$U_T^{(1,n)}(\mu, K_1) + U_T^{(2,n)}(\mu, K_2).$$

The Result We show that an asymptotically optimal sequence can be constructed by simply employing $\mu^{*,F}$ in every pre-limit station. This turns out to be a consequence of an elegant FSLLN-type result (see Proposition 1.2.4). The shortcomings (possible "waste" of service in the first station or downward pushing in the second one) of this service discipline in the pre-limit can be easily rectified using a bit more general form of the $o(\frac{1}{\sqrt{n}})$-rule from the result for the pre-limit sequences in the single station setting. However, we have already gathered understanding of the consequences of imposing the $o(\frac{1}{\sqrt{n}})$-rule on a fluid-optimal policy when we considered the single station.

2.3.2 The Infinite Buffer Context

We retain the arrival and departure processes described in the finite-buffer setting and set the buffers in both stations to be infinite. The constants K_1 and K_2 are now understood as thresholds. There is a unit penalty accumulated in each station for the time the queue length spends above its respective threshold.

We immediately delve into the sequence of uniformly accelerated systems. The queue length processes are

$$\begin{aligned} Q^{(1,n)}(\mu) &= N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu)) + L^{(1,n)}(\mu), \\ Q^{(2,n)}(\mu) &= N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu) - N_2^-(n\mathcal{I}(\mu^2)) + L^{(2,n)}(\mu), \end{aligned} \tag{2.3.1}$$

where the processes $L^{(1,n)}(\mu)$ and $L^{(2,n)}(\mu)$ are the regulating terms arising from the application of the one-dimensional reflection map Γ of Definition 1.2.1 to processes $N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu))$ and $N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu) - N_2^-(n\mathcal{I}(\mu^2))$, respectively. One system in this sequence is depicted in Figure 2.8.

The Fluid Limit [***Fully covered in Section 4.4***]

The Model Given a fixed deterministic service discipline μ, we consider the following pair of queue lengths in the fluid-limit

$$\bar{Q}^{(1)}(\mu) = \mathcal{I}(\lambda) - \mathcal{I}(\mu) + \bar{L}^{(1)}(\mu),$$
$$\bar{Q}^{(2)}(\mu) = \mathcal{I}(\mu) - \bar{L}^{(1)}(\mu) - \mathcal{I}(\mu^2) + \bar{L}^{(2)}(\mu),$$

where

$$\bar{L}^{(1)}(\mu) = \sup_{s\leq\cdot}[-\mathcal{I}_s(\lambda - \mu)],$$
$$\bar{L}^{(2)}(\mu) = \sup_{s\leq\cdot}[-\mathcal{I}_s(\mu) + \bar{L}^{(1)}_s(\mu) + \mathcal{I}_s(\mu^2)].$$

The Performance Measure In analogy with the single station case, the constants K_1 and K_2 are considered as thresholds and the fluid-limit performance measure is given by

$$\text{meas}\{t \in [0,T] : \bar{Q}^{(1)}_t(\mu) > K_1\} + \text{meas}\{t \in [0,T] : \bar{Q}^{(2)}_t(\mu) > K_2\},$$

for every admissible deterministic control μ.

The Result Again a comparison with the pooled counterpart (shown in Figure 2.9) delivers a class of fluid-optimal disciplines defined by the following three features:

- there is no upward pushing in the first station,
- at no time does the queue length in the second station cross over the threshold K_2,
- at any instant the first queue is above its threshold only if the pooled queue is above its threshold.

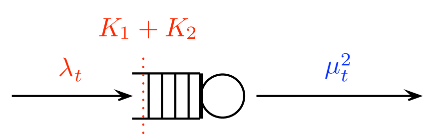

Figure 2.9: The pooled queue - Infinite buffer

The particular fluid-optimal discipline we identify in the present case follows the same rules as the fluid-optimal policy $\mu^{*,F}$ of the finite-buffer case: no lower regulation in the first station and no penalty incurred in the second station. The implemented disciplines can have different explicit forms, as the reflection maps are different in the infinite and the finite buffer cases.

The Pre-limit Sequence [*Fully covered in Section 4.6*]

The Model The lengths of queues in the pre-limit systems are already displayed in (2.3.1).

The Performance Measure For every n and every admissible control μ, the performance measure is

$$\int_0^T \left[\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,n)}(\mu)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,n)}(\mu)>K_2\}} \right] dt.$$

We wish to asymptotically minimize its expected value.

The Result We propose an asymptotically optimal class of sequences of stochastic controls. This class consists of sequences of service rates that depend on the state of the system - the rates are set to zero whenever the first queue is empty or the second queue is at the upper threshold. In all other states of the systems, the controlled rates are set to increase with the index n "much faster" than the uniformly accelerated rates of arrivals in the first station and service in the second. The exact nature of the asymptotic behavior within the proposed class of controls is elaborated on in Section 4.6. Moreover, we show the existence of a set of parameters for which there is no deterministic asymptotically optimal control sequence.

2.4 The Overall Strategy

To complete the overview of the results, we exhibit the "finite automaton" diagram in Figure 2.10 which contains the relationships between the different control problems in the single station case. Starting with the upper left corner and following the arrows, we develop the train of thought which is the backdrop for all the control problems tackled individually later in the text. A step-by-step verbal description would go as follows:

1. We start with the general formulation of the control problem aiming to minimize the number of lost jobs.

2. Realizing the potential complications arising from the use of the two-sided regulation mapping, we decide to focus on a simpler network - the one with an infinite buffer - and consider the control problem associated with the time the queue length spends across a given threshold. This is not the same problem as the original one, but exhibits similarities.

3. The exact analysis is hard, so we proceed with a commonly used strategy of embedding the "real" system in a sequence of uniformly accelerated ones.

4. We pose and solve the associated control problems in the fluid-limit setting, first for the infinite buffer and then the finite buffer case.

Figure 2.10: The Strategy - Single station

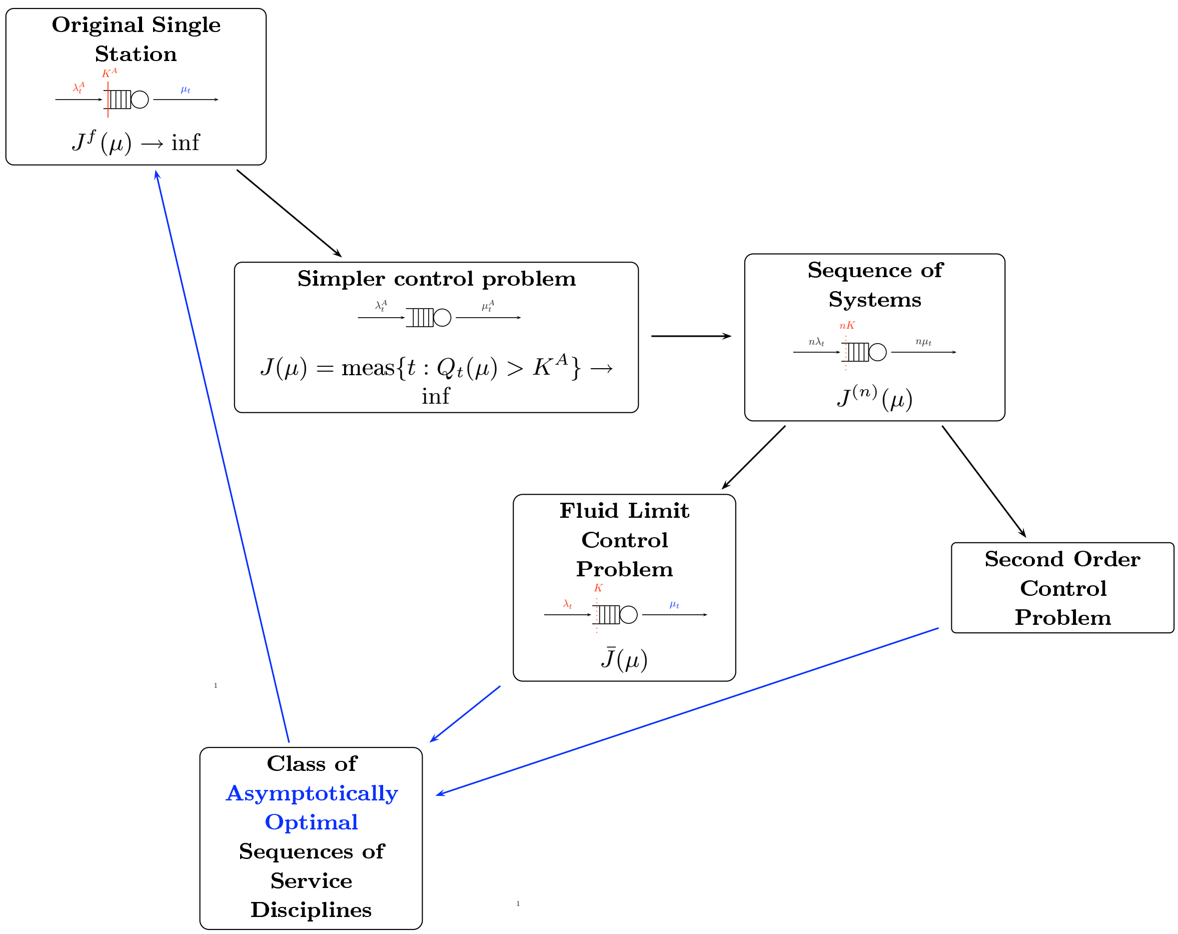

Figure 2.11: A fluid-optimal policy which performs poorly in the pre-limit.

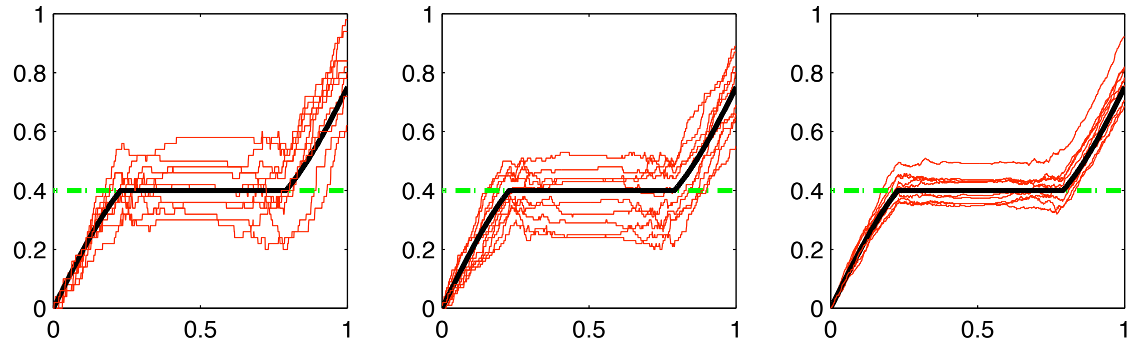

5. We introduce the novel notion of second order optimality, formulate the resulting control problem and provide an optimal class of policies. This is conducted only in the infinite buffer setting.

6. Informed by the optimal classes in the fluid and second order settings, we provide an asymptotically optimal class of policies for the pre-limit sequence of systems - both for finite and infinite buffers.

Let us briefly return to the questions we posed in the end of Subsection 2.1.1.

(i) The formulation of the second-order optimal control problem and the rationale behind it are described above in Subsection 2.2.2.

(ii) The immediacy and self-containment of the asymptotic control problems in the sequel may prompt the reader to question the rationale behind venturing into the second order control problem at all. The reasons for this consideration are two-fold. First, considering the second order problems separately gave insights in the relationship between the asymptotic control problems featuring the expected value of the penalty and the control problems employing asymptotic optimality in an almost sure sense. For more on this subject, see Appendix C.5. Second, there exist settings in which the fluid-limit analysis does not give immediate direction to how an asymptotically optimal sequence is constructed and considering the second order control problem enables the controller to obtain insight into the pre-limit behavior of the system (see Section 3.3). The illustration of this phenomenon is given in Figure 2.11. In this example, it is clear that the controller would try to keep the mean queue length sufficiently away from the boundary. However, it is not *a priori* clear what "sufficiently away" stands for. The second-order optimal control problems helps us to answer this question.

(iii), (iv) In their full generality, these two question remain an ambition for future research. In the context of asymptotically optimal control problems considered in this thesis, we find that all but one (the one including the tandem system with infinite buffers, as described above in Subsection 2.3.2) allow for deterministic asymptotically optimal policies. Also, in all the problems we look at in this thesis, we provide classes of asymptotically optimal policies whose limiting performance does not depend on the choice of the index corresponding to the actual system in the uniform acceleration scheme.

2.5 On Simplifications and Extensions

All of the simplifications we state below have the purpose of not adding extra burden to the notation, and have no methodological consequences whatsoever. We now list three simplifying assumptions that were made in the descriptions of the optimal control problems above. These assumptions were made purely for ease of exposition and notational convenience - a slight modification of the proofs in the thesis would allow one to easily relax these assumptions.

First, throughout this work, all queues are assumed to be initially empty. This is nothing but a convenience - all results can be carried over quite immediately to the case of nonzero initial conditions. Secondly, the single-station study can be extended straightforwardly to the system with feedback, i.e., the system in which some of the served jobs return to the back of the queue with a given (not necessarily constant) probability. Finally, the same remark applies to the addition of a feedback loop between the second station and the first queue of the tandem system.

Chapter 3

Single Station

3.1 The Control Problem Set-up

3.1.1 The Model

As outlined in Subsection 2.2.1, we focus on a single station with one server serving one job class, where λ denotes the arrival rate and μ the service rate. The rate λ is assumed deterministic, nonnegative and integrable. However, μ can be random, taking values in $\mathbb{L}^1_+[0,T]$. Reasons for permitting stochastic service disciplines, as well as a description of their particular structure, are provided in Subsection 3.1.3. The randomness in the arrival and the potential service processes is modeled by independent unit Poisson processes N^+ and N^-, time-changed by the rates λ and μ, respectively.

The system we look at has a finite buffer capacity denoted by K, which may cause loss of jobs from the system. More precisely, whenever the buffer is full, a new incoming job cannot queue up and leaves the system without being served.

3.1.2 A Sequence of Systems

As announced in Subsection 2.2.1, we proceed to construct and study a sequence of uniformly accelerated systems based on the above data and modeling assumptions. For each n and for a given admissible service discipline μ, the netput process is

$$X^{(n)}(\mu) = N^+(n\mathcal{I}(\lambda)) - N^-(n\mathcal{I}(\mu)). \tag{3.1.1}$$

Applying the two-sided reflection maps Γ^{nK} of Definition 1.2.2 to $X^{(n)}(\mu)$, we obtain the lengths of the queues[1]

$$Q^{(n)}(\mu) = \Gamma^{nK}(X^{(n)}(\mu)). \tag{3.1.2}$$

3.1.3 Space of Well-Adapted Service Disciplines

We will shortly face the task of finding service disciplines that are optimal for a certain performance measure. We wish to be at liberty to choose a suitable instantaneous service rate μ_t, for all $t \in [0, T]$, with the assumption that the whole past of the system is known at any time. This, in effect, means we must allow for non-deterministic μ, when controlling the prelimit processes $Q^{(n)}$. Furthermore, it is expected of the controller of the system to be aware of certain properties of the arrival process. Namely, statistics of the past behavior of the system to be controlled (or other systems akin to it) can give a reasonable estimate of the average rate of arrivals as a function of time. We will, hence, assume that $\lambda \in \mathbb{L}^1_+[0, T]$ is a completely known (deterministic) function.

The assumptions just described - both on the model of the system given in terms of Poisson processes, and the flow of information containing the entire arrival rate and the past of the system - are natural. However, the formal mathematical description of these concepts turns out to be quite technical. The definition of the space accommodating the service disciplines is given separately in Appendix B.2. Let us rely on that construction and from now on refer to the set of all admissible service disciplines in the n^{th} system as $\mathcal{L}^{(n)}$, for all $n \in \mathbb{N}$. The space of all sequences $\{\mu_n\}$ such that $\mu_n \in \mathcal{L}^{(n)}$, for every $n \in \mathbb{N}$, will be denoted by $\mathcal{L}$.

The specific constraint we impose on the admissible service disciplines addresses the total amount of service available on the interval $[0, T]$. Formally, given a constant $m \in \mathbb{R} \cup \{\infty\}$, our attention will be restricted, for any n, to $\mu \in \mathcal{L}^{(n)}$ such that $\mathcal{I}_T(\mu) \leq m$, almost surely. The space of controls satisfying these conditions will be denoted by $\mathcal{L}^{(n)}(m)$, and the space of all admissible sequences conforming to the extra constraint will be called $\mathcal{L}(m)$. Note that when $\mu \in \mathcal{L}^{(n)}$, the control effort expended in the n^{th} system, namely $n\mathcal{I}_T(\mu)$ (see (3.1.1)), is bounded from above by nm.

3.1.4 Sequence of Performance Measures

Given the system described above, it is of interest to try to minimize the number of jobs lost due to the finite buffer capacity. Our first step is to formally express this quantity as a function of the admissible service discipline used. Thanks to the decomposition (1.2.5) of Definition 1.2.2,

[1] We omit the explicit mention of the threshold K from the notation for the queue length for brevity's sake. As there is no chance of confusion, the same convention is applied in later considerations of queues with finite buffer capacities.

we can expand the queue length in the n^{th} system as

$$Q^{(n)}(\mu) = X^{(n)}(\mu) + L^{(n)}(\mu) - U^{(n)}(\mu), \text{ for any } \mu \in \mathcal{L}^{(n)}, \tag{3.1.3}$$

where the random processes $L^{(n)}(\mu)$ and $U^{(n)}(\mu)$ are associated with $X^{(n)}(\mu)$ and nK in the sense of Definition 1.2.2. Now, we can explicitly define the corresponding performance measure $J_F^{(n)} : \mathcal{L}^{(n)}(m) \to \mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ as

$$J_F^{(n)}(\mu) = \frac{1}{n} U_T^{(n)}(\mu), \text{ for every } \mu \in \mathcal{L}^{(n)}(m). \tag{3.1.4}$$

This performance measure is reasonably difficult to analyze, and so we decide to consider a simpler control problem closely related to the one proposed above. We elaborate on this in the following section.

3.1.5 The Infinite Buffer Problem

Instead of the system we discussed in Section 3.1 (which was depicted in Figure 2.2), we now focus on the system with infinite buffer capacity but with the same exogenous arrival process. The uniform acceleration scheme delivers a sequence of systems with queue lengths

$$Q^{(n)}(\mu) = \Gamma(X^{(n)}(\mu)), \text{ for every } n, \tag{3.1.5}$$

where $X^{(n)}(\mu)$ is the netput process of (3.1.1) and Γ is the one-sided reflection map introduced in Definition 1.2.1. As stated in Subsection 2.2.2, in the present context the positive constant K plays the role of a given threshold. For every n, we define the mapping $J^{(n)} : \mathcal{L}^{(n)}(m) \longrightarrow \mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ as

$$J^{(n)}(\mu) = \int_0^T \mathbf{1}_{\{\frac{1}{n} Q_t^{(n)}(\mu) > K\}} dt, \text{ for every } \mu \in \mathcal{L}^{(n)}(m). \tag{3.1.6}$$

The performance measures (3.1.6) can be understood as aggregated unit penalties incurred when the (normalized) queues exceed the threshold K. For instance, this can occur in a manufacturing chain due to a fixed cost of an additional storage facility needed in case the capacity K of the present facility is exceeded. Our goal is to first produce a sequence of admissible service disciplines minimizing $J^{(n)}$ in the asymptotic sense which will be elaborated on in Section 3.4.

3.2 Fluid Limit

3.2.1 The Infinite-Buffer Case

We begin the analysis by approximating the sequence of systems in (3.1.5) by the fluid limit system. For any given service discipline $\mu \in \mathbb{L}_+^1[0, T]$, the Functional Strong Law of Large

Numbers applied to the sequence of random processes $\{X^{(n)}\}$ defined in (3.1.1), yields

$$\frac{1}{n}X^{(n)}(\mu) \longrightarrow \bar{X}(\mu),\ a.s. \tag{3.2.1}$$

where

$$\bar{X}(\mu) = \mathcal{I}(\lambda - \mu). \tag{3.2.2}$$

The convergence of processes in (3.2.1) is with respect to the uniform topology on the space $\mathcal{D}$. The continuity of the reflection map Γ then gives us

$$\frac{1}{n}Q^{(n)}(\mu) \longrightarrow \bar{Q}(\mu), \tag{3.2.3}$$

also in the uniform topology on $\mathcal{D}$, where

$$\bar{Q}(\mu) = \Gamma(\bar{X}(\mu)). \tag{3.2.4}$$

The expression (1.2.2) provides a decomposition of $\bar{Q}(\mu)$ of the form

$$\bar{Q}(\mu) = \bar{X}(\mu) + \bar{L}(\mu),$$

with $\bar{L}(\mu) = \sup_{s\leq\cdot}[-\bar{X}(\mu)]^{+}$. The natural analogue of the sequence of performance measures in (3.1.6) in the fluid-limit setting is the performance measure $\bar{J} : \mathbb{L}^1_+[0,T] \to \mathbb{R}_+$, defined by

$$\bar{J}(\mu) = \int_0^T \mathbf{1}_{\{\bar{Q}_t(\mu)>K\}}dt. \tag{3.2.5}$$

In the context of the deterministic fluid approximation in (3.2.3), we choose to consider only deterministic μ. Also, the restriction $\mathcal{I}_T(\mu) \leq m$ on the cumulative amount of service is inherited from the original problem described in Section 3.1. Hence, the space of admissible disciplines for the minimization problem associated with the performance measure $\bar{J}$ is

$$\bar{\mathcal{L}}(m) = \{\mu \in \mathbb{L}^1_+[0,T] : \mathcal{I}(\mu)_T \leq m\}. \tag{3.2.6}$$

The simplicity of the stated control problem allows more general results which are exhibited in Appendix C.1. To ensure a smoother flow of the argument certain statements are included both here and in the appendix. First, mappings generating the fluid netput and queue length processes are formally defined.

Definition 3.2.1. Mappings $x, q : \mathbb{R}_+ \times \mathbb{L}^1_+[0,T] \times \mathbb{L}^1_+[0,T] \to \mathcal{C}$ are defined as

$$x(q_0, \lambda, \mu) = q_0 + \mathcal{I}(\lambda - \mu) \quad \text{and} \quad q = \Gamma \circ x,$$

for $q_0 \geq 0$ and where λ and μ are nonnegative, integrable functions on the segment $[0,T]$. We say that the process $q : [0,T] \to \mathbb{R}$ is the fluid-limit queue length process *generated* by λ as arrival rate and μ as service rate, starting at q_0.

This definition allows us to rewrite the processes in (3.2.2) and (3.2.4) as $\bar{X}(\mu) = x(0, \lambda, \mu)$ and $\bar{Q}(\mu) = q(0, \lambda, \mu)$. Next, we provide a particular partial ordering on the space of all possible queue lengths.

Definition 3.2.2. Let $q_1, q_2 \in \mathcal{C}$ be such that $q_1 \leq q_2$ almost everywhere. Then we say that q_2 *dominates* q_1, and write $q_1 \preceq q_2$.

The stage is set for the main definition.

Definition 3.2.3. The mapping $j : \mathbb{R}_+ \times \mathbb{L}^1_+[0,T] \times \mathbb{L}^1_+[0,T] \to \mathbb{R}$ is called an *increasing* (resp.*decreasing*) performance measure, if for all $(q_0^i, \lambda^i, \mu^i), i = 1, 2$, such that $q(q_0^1, \lambda^1, \mu^1) \preceq q(q_0^2, \lambda^2, \mu^2)$ (resp. $q(q_0^1, \lambda^1, \mu^1) \succeq q(q_0^2, \lambda^2, \mu^2)$) we have $j(q_0^1, \lambda^1, \mu^1) \leq j(q_0^2, \lambda^2, \mu^2)$ (resp. $j(q_0^1, \lambda^1, \mu^1) \geq j(q_0^2, \lambda^2, \mu^2)$). If j is either decreasing or increasing, it will be referred to as a *monotone* performance measure.

Let us continue by introducing τ as the (right-continuous) inverse of the fluid limit $\mathcal{I}(\lambda)$ of the arrival process, i.e., for every $l \in \mathbb{R}_+$, we let

$$\tau(l) = \inf\{t \geq 0 : \mathcal{I}_t(\lambda) > l\} \wedge T. \tag{3.2.7}$$

Lemma 3.2.4. *The mapping $\bar{J}$ defined in (3.2.5) is an increasing performance measure, in the sense that for all $\mu_1, \mu_2 \in \bar{\mathcal{L}}(m)$*

$$\bar{Q}(\mu_1) \preceq \bar{Q}(\mu_2) \Rightarrow \bar{J}(\mu_1) \leq \bar{J}(\mu_2).$$

The optimal (minimal) value that $\bar{J}$ attains on $\bar{\mathcal{L}}(m)$ is

$$J^* = T - \tau(K + m), \tag{3.2.8}$$

where the mapping τ is given in (3.2.7).

Proof. Let us commence with $\mu_1, \mu_2 \in \bar{\mathcal{L}}(m)$ such that $\bar{Q}(\mu_1) \preceq \bar{Q}(\mu_2)$. Then, by definition, $\bar{Q}_t(\mu_1) \leq \bar{Q}_t(\mu_2)$ for almost all $t \in [0,T]$, implying that for almost all t such that $\bar{Q}_t(\mu_1) > K$, it must be that $\bar{Q}_t(\mu_2) > K$. Integrating over $[0,T]$, we obtain the inequality $\bar{J}(\mu_1) \leq \bar{J}(\mu_2)$.

By Proposition C.1.8, the infimum of the mapping $\bar{J}$ on $\bar{\mathcal{L}}(m)$ is attained at $\mu^* = \lambda \mathbf{1}_{[0,\tau(m)]}$. Hence, it suffices to evaluate $\bar{J}(\mu^*)$. For all $t \leq \tau(m)$, the resulting queue length is identically zero, and for $t > \tau(m)$, $\bar{Q}_t(\mu^*) = \mathcal{I}_t(\lambda) - m$. Thus

$$\bar{J}(\mu^*) = \text{meas}\{t > \tau(m) : \mathcal{I}_t(\lambda) - m > K\} = T - \tau(K + m).$$

□

Finally, we state a criterion for optimality with respect to the performance measure $\bar{J}$ for the case in which the restriction on the cumulative service available has nontrivial consequences on the possible penalty.

Theorem 3.2.5. *Assume that $\mathcal{I}_T(\lambda) > K + m$. Then for $\mu \in \bar{\mathcal{L}}(m)$, $\bar{J}(\mu) = J^*$ iff $\bar{L}_T(\mu) = 0$ and $\mathcal{I}_t(\mu) \geq (\mathcal{I}_t(\lambda) - K)^+$ for $t < \tau(K+m)$.*

Proof. First note that for any admissible $\mu \in \bar{\mathcal{L}}(m)$, for all $t > \tau(K+m)$ we have

$$\bar{Q}_t(\mu) \geq \bar{X}_t(\mu) = \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu) > K + m - m = K. \tag{3.2.9}$$

Thus, by the complementarity condition (ii) of Definition 1.2.1, we conclude that $\bar{L}$ is "flat" on the interval $(\tau(K+m), T]$, and so $\bar{L}_T(\mu) = \bar{L}_{\tau(K+m)}(\mu)$.

Now, let $\mu \in \bar{\mathcal{L}}(m)$ be optimal for the performance measure $\bar{J}$. Contrary to the claim, let us assume $\bar{L}_T(\mu) > 0$. Then $\bar{L}_{\tau(K+m)}(\mu) = \bar{L}_T(\mu) > 0$. So, already at $\tau(K+m)$ the queue length strictly exceeds the threshold K, as is readily seen from

$$\bar{Q}_{\tau(K+m)}(\mu) = \mathcal{I}_{\tau(K+m)}(\lambda) - \mathcal{I}_{\tau(K+m)}(\mu) + \bar{L}_{\tau(K+m)}(\mu) > K + m - m = K, \tag{3.2.10}$$

where we have used the fact that $\mathcal{I}_{\tau(K+m)}(\mu) \leq m$, since $\mu \in \bar{\mathcal{L}}(m)$. Let t_o be the last point before $\tau(K+m)$ at which there is no penalty, i.e.,

$$t_o := \sup\{s \leq \tau(K+m) : \bar{Q}_s(\mu) \leq K\}. \tag{3.2.11}$$

By (3.2.10) and the continuity of $\bar{Q}(\mu)$, we have $t_o < \tau(K+m)$. This observation, when combined with (3.2.9), implies that $\bar{J}(\mu) \geq T - t_o > J^*$, which leads to a contradiction.

Let us now assume that for an optimal μ there exists an instant $t_p < \tau(K+m)$ such that $\mathcal{I}_{t_p}(\mu) < (\mathcal{I}_{t_p}(\lambda) - K)^+$. Then we necessarily have $t_p > \tau(K)$ and $(\mathcal{I}_{t_p}(\lambda) - K)^+ = \mathcal{I}_{t_p}(\lambda) - K$. Hence, $\bar{Q}_{t_p}(\mu) \geq \bar{X}_{t_p}(\mu) = \mathcal{I}_{t_p}(\lambda) - \mathcal{I}_{t_p}(\mu) > K$. By the continuity of $\bar{Q}(\mu)$, there is an interval (t_l, t_r) containing t_p such that $\bar{Q}_t(\mu) > K$ for all $t \in (t_l, t_r)$. By the same reasoning as in (3.2.9), we have

$$\text{meas}\{t \in [\tau(K+m), T] : \bar{Q}_t(\mu) > K\} = T - \tau(K+m).$$

Hence, $J(\mu) \geq (T - \tau(K+m)) + (t_r - t_l) > J^*$. This, again, results in a contradiction. Together with the last paragraph, this proves the "only if" part of the theorem.

As for the easier direction in the equivalence in the theorem, let $\mu \in \bar{\mathcal{L}}(m)$ satisfy $L_T(\mu) = 0$ and $\mathcal{I}_t(\mu) \geq (\mathcal{I}_t(\lambda) - K)^+$ for $t < \tau(K+m)$. Then, for all $t \in [0, \tau(K+m)]$

$$\bar{Q}_t(\mu) = \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu) \leq \mathcal{I}_t(\lambda) - (\mathcal{I}_t(\lambda) - K)^+ \leq K.$$

By (3.2.9) again, $\bar{J}(\mu) = T - \tau(K+m) = J^*$. ∎

3.2.2 Consequences in the Finite Buffer Setting

In this section, $\{X^{(n)}(\mu)\}$ will continue to represent the sequence of netput processes defined in (3.1.1) and $\{Q^{(n)}(\mu)\}$ the sequence of queue length processes (depending on the service discipline μ). Recall that in the finite buffer context they satisfy the relation $Q^{(n)}(\mu) = \Gamma^{nK}(X^{(n)}(\mu))$. Since this subsection deals entirely with the finite buffer case, this should not be confused with the previous subsection where $Q^{(n)}(\mu)$ is defined in (3.1.5). Proposition 1.2.4 and the convergence result (3.2.1) allow us to conclude that

$$\frac{1}{n}Q^{(n)}(\mu) \to \bar{Q}(\mu) = \Gamma^K(\bar{X}(\mu)), \text{ for every } \mu \in \mathbb{L}^1_+[0,T],$$

where $\bar{X}(\mu)$ is defined in (3.2.2). In the already familiar fashion, we write

$$\bar{Q}(\mu) = \bar{X}(\mu) + \bar{L}(\mu) - \bar{U}(\mu), \tag{3.2.12}$$

where $\bar{L}(\mu)$ and $\bar{U}(\mu)$ are associated with $\bar{X}(\mu)$ and K in the sense of Definition 1.2.2. The fluid-limit version of the performance measures (3.1.4) is

$$\bar{J}^F(\mu) = \bar{U}_T(\mu), \text{ for } \mu \in \bar{\mathcal{L}}(m),$$

with $\bar{\mathcal{L}}(m)$ defined in (3.2.6).

Next, we determine a (non-trivial) lower bound for this performance measure on the space $\bar{\mathcal{L}}(m)$. This choice of a lower bound is inspired by the results gained in the infinite-buffer setting. More precisely, we concentrate on the service disciplines recognized as fluid-optimal in the context of Theorem 3.2.5 to guess the lower bound in the following lemma.

Lemma 3.2.6. *For all $\mu \in \bar{\mathcal{L}}(m)$, we have*

$$\bar{J}^F(\mu) \geq (\mathcal{I}_T(\lambda) - K - m)^+. \tag{3.2.13}$$

Moreover, the minimum of the mapping $\bar{J}^F$ over all $\mu \in \bar{\mathcal{L}}(m)$ is achieved by the policy $\mu^* = \lambda \mathbf{1}_{[0,\tau(m)]}$.

Proof. Equation (3.2.12) and the fact that $\bar{L}(\mu)$ is by definition nonnegative lead us to conclude that

$$\bar{J}^F(\mu) = \bar{U}_T(\mu) \geq \mathcal{I}_T(\lambda) - \mathcal{I}_T(\mu) - \bar{Q}_T(\mu).$$

Since $\mu \in \bar{\mathcal{L}}(m)$ and $\bar{Q}_T(\mu) \leq K$, the above inequality implies

$$\bar{J}^F(\mu) = \bar{U}_T(\mu) \geq \mathcal{I}_T(\lambda) - m - K.$$

Since, by definition, $\bar{J}^F$ is nonnegative, this completes the proof of (3.2.13).

Now we verify that the lower bound in (3.2.13) is indeed attained at μ^*. Thanks to a well-known identity connecting the upper and lower regulators in the two-sided reflection map (see, e.g., equation (1.9) in [KLRS06] or Section 2.3 of [Har90]), we have

$$\bar{J}^F(\mu^*) = U_T(\mu^*) = \sup_{s\le T}[\bar{X}_s(\mu^*) + \bar{L}_s(\mu^*) - K]^+.$$

Since $\bar{L}(\mu^*) \equiv 0$ and $\mathcal{I}(\mu^*) = \mathcal{I}(\lambda) \wedge m$, the last equality implies

$$\bar{J}^F(\mu^*) = \sup_{s\le T}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\lambda) \wedge m - K]^+ = \sup_{s\le T}[(\mathcal{I}_s(\lambda) - m)^+ - K]^+ = (\mathcal{I}_T(\lambda) - m - K)^+.$$

□

The following analogue of Theorem 3.2.5 can be established using the same arguments as in the proof of Theorem 3.2.5.

Theorem 3.2.7. *Assume that $\mathcal{I}_T(\lambda) > K + m$. Then for $\mu \in \bar{\mathcal{L}}(m)$, $\bar{J}^F(\mu) = (\mathcal{I}_T(\lambda) - m - K)^+$ iff $\bar{L}_T(\mu) = 0$ and $\mathcal{I}_t(\mu) \ge (\mathcal{I}_t(\lambda) - K)^+$ for $t < \tau(K + m)$.*

3.3 Second Order Approximation

In this section we still restrict our attention to deterministic service disciplines, i.e., if it is not explicitly noted otherwise, whenever we refer to a service discipline we will have a deterministic one in mind.

3.3.1 Some Important Processes

We start by recalling Definition 3.2.1, and proceed to define for any service discipline $\mu \in \mathbb{L}^1_+[0,T]$ the deterministic processes

$$\bar{X}(\mu) = x(0,\lambda,\mu) \qquad \text{and} \qquad \bar{Q}(\mu) = q(0,\lambda,\mu), \tag{3.3.1}$$

for a fixed $\lambda \in \mathbb{L}^1_+[0,T]$. We continue to interpret λ as a known, deterministic arrival rate and consider it fixed for the rest of this section.

For any μ, the process $\bar{X}(\mu)$ introduced in (3.3.1) induces a set-valued process $\Phi_{-\bar{X}(\mu)}$ defined by

$$\Phi_{-\bar{X}(\mu)}(t) = \{s \le t : -\bar{X}_s(\mu) = \sup_{u\le t}[-\bar{X}_u(\mu)]\} \text{ for all } t \in [0,T]. \tag{3.3.2}$$

For some useful facts about $\Phi_{-\bar{X}(\mu)}$, the reader is directed to Section 9.3 in [Whi02a].

Throughout the rest of this section W will stand for a standard Brownian motion on the (possibly enlarged) probability space $(\Omega, \mathcal{F}, \mathbb{P})$. For any $\mu \in \mathbb{L}^1_+[0,T]$, using (3.3.2), we define the random process

$$\hat{Q}(\mu) = W(\mathcal{I}(\lambda + \mu)) + \sup_{s \in \Phi_{-\bar{X}(\mu)}(\cdot)} [-W(\mathcal{I}_s(\lambda + \mu))]. \tag{3.3.3}$$

This process arises in the second-order approximation (for more details, see Theorem 1.2.7 and Subsection 2.2.2).

3.3.2 A Sequence of Control Problems

We now consider a sequence of performance measures defined by the mappings $\hat{J}^{(n)} : \mathbb{L}^1_+[0,T] \to \mathbb{L}^0_+(\Omega, \mathcal{F}, \mathbb{P})$, which are given by

$$\hat{J}^{(n)}(\mu) = \text{meas}\left\{ t \in [0,T] : \bar{Q}_t(\mu) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu) > K \right\}, \tag{3.3.4}$$

with $\bar{Q}(\mu)$ and $\hat{Q}(\mu)$ as above. As was discussed in Subsection 2.2.2, the processes

$$\bar{Q}_t(\mu) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu)$$

featured in the definition of the sequence of mappings $\{\hat{J}^{(n)}\}$ above provide a second-order approximation for the sequence $\{\frac{1}{n}Q^{(n)}(\mu)\}$. We are not interested in minimizing the value of $\hat{J}^{(n)}$ for any fixed index n. What we wish to explore is the asymptotic behavior of the sequence of performance measures as $n \to \infty$. Hence, it is sensible to have the following definition.

Definition 3.3.1. A sequence $\{\mu_n^*\}$ in $\bar{\mathcal{L}}(m)$ is called *second order optimal* for the sequence of performance measures $\{\hat{J}^{(n)}\}$ if

$$\liminf_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n) - \hat{J}^{(n)}(\mu_n^*)] \geq 0$$

for any other sequence $\{\mu_n\}$ of deterministic admissible service disciplines.

The following lemma contains a more operational criterion for determining second order optimality.

Lemma 3.3.2. *Let $\hat{J}^*$ be a constant such that for all sequences $\{\mu_n\}$ in $\bar{\mathcal{L}}(m)$ we have*

$$\liminf_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n) - \hat{J}^*] \geq 0,$$

and suppose that $\{\mu_n^*\}$ is a sequence in $\bar{\mathcal{L}}(m)$, such that

$$\lim_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n^*)] = \hat{J}^*. \tag{3.3.5}$$

Then $\{\mu_n^*\}$ is second order optimal for the sequence $\{\hat{J}^{(n)}\}$.

Proof. Let $\hat{J}^*$ and $\{\mu_n^*\}$ satisfy the assumptions of the lemma. Then for all deterministic admissible sequences $\{\mu_n\}$,

$$\begin{aligned}\liminf_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n) - \hat{J}^{(n)}(\mu_n^*)] &\geq \liminf_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n) - \hat{J}^*] + \liminf_{n\to\infty} \mathbb{E}[\hat{J}^* - \hat{J}^{(n)}(\mu_n^*)] \\ &\geq \lim_{n\to\infty} \mathbb{E}[\hat{J}^* - \hat{J}^{(n)}(\mu_n^*)] = 0.\end{aligned}$$

□

The previous lemma allows us to couple the direct search for the sequence $\{\mu_n^*\}$ with the search for a suitable asymptotic lower bound $\hat{J}^*$. At first sight, the benefits of this intermediate step can easily be challenged. Clearly, the success of the procedure relies heavily on a clever choice of $\hat{J}^*$. However, we come to this apparent obstacle armed with the intuition gained in the fluid-limit analysis. This connection is what makes the two-step procedure more natural. The details of how $\hat{J}^*$ is determined are presented in Appendix C.3.

Proposition 3.3.3. *Let the constant* $\hat{J}^*$ *be defined as*

$$\hat{J}^* = \frac{1}{2}(\tau(K+m) - \tau((K+m)-)) + T - \tau(K+m). \tag{3.3.6}$$

Then for all deterministic admissible sequences $\{\mu_n\}$,

$$\liminf_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\mu_n) - \hat{J}^*] \geq 0.$$

Proof. This is a direct consequence of Lemma C.3.2 and Corollary C.3.4. □

The first part of the criterion from Lemma 3.3.2 is, thus, satisfied by $\hat{J}^*$ in (3.3.6). What remains to be done, in order to fully apply Lemma 3.3.2, is to propose a particular admissible sequence of service disciplines whose limiting penalty will match that of $\hat{J}^*$, as specified in (3.3.5). The next subsection is dedicated to this task.

3.3.3 A Second Order Optimal Class

Assumption 3.3.4. Let $\{\eta_n\}$ be a sequence in $(0,1)$ satisfying

$$n(1-\eta_n)^2 \longrightarrow \infty, \text{ as } n \to \infty.$$

Furthermore, let $\{\hat{\mu}_n\}$ be a sequence of integrable nonnegative functions on $[0,T]$, defined as

$$\hat{\mu}_n = \lambda \mathbf{1}_{[\tau(\eta_n K), \tau(\eta_n K + m)]}. \tag{3.3.7}$$

In words, the service discipline $\hat{\mu}_n$ "does nothing" until the cumulative arrival rate reaches the level $\eta_n K$, at which moment it starts matching the arrival rate, and does so until it "runs out of fuel". The cumulative service disciplines can be written as $\mathcal{I}_t(\hat{\mu}_n) = (\mathcal{I}_t(\lambda) - \eta_n K)^+ \wedge m$. It is easily seen that $\{\hat{\mu}_n\}$ is a sequence of admissible disciplines.

For any n, the process $\bar{Q}(\hat{\mu}_n)$ defined in (3.2.4) has the form

$$\bar{Q}_t(\hat{\mu}_n) = \begin{cases} \mathcal{I}_t(\lambda) & \text{for } 0 \le t \le \tau(\eta_n K), \\ \eta_n K & \text{for } \tau(\eta_n K) \le t \le \tau(\eta_n K + m), \\ \mathcal{I}_t(\lambda) - m & \text{for } \tau(\eta_n K + m) \le t \le T. \end{cases} \tag{3.3.8}$$

This observation has the following simple consequence.

Lemma 3.3.5. *For every n, the service discipline $\hat{\mu}_n$ is fluid-optimal. Equivalently,*

$$\bar{J}(\hat{\mu}_n) = T - \tau(K + m) \, \textit{for every } n.$$

Another helpful feature of the sequence $\{\hat{\mu}_n\}$ is the simplification of the form of the processes $\Phi_{-\bar{X}(\hat{\mu}_n)}$ defined in (3.3.2), under a mild assumption that $\lambda > 0$ on some neighborhood of zero. In fact, for every n we simply get

$$\Phi_{-\bar{X}(\hat{\mu}_n)}(t) = \{s \le t : -\bar{X}_s(\hat{\mu}_n) = \sup_{u \le t}[-\bar{X}_u(\hat{\mu}_n)] = 0\} = \{0\}, \text{ for every } t.$$

This fact enables us to display the stochastic processes $\hat{Q}(\hat{\mu}_n)$ introduced in (3.3.3) as

$$\hat{Q}(\hat{\mu}_n) = W(\mathcal{I}(\lambda + \hat{\mu}_n)). \tag{3.3.9}$$

In words, the processes $\hat{Q}(\hat{\mu}_n)$ become merely Brownian motions with different time-changes.

Theorem 3.3.6. *The sequence $\{\hat{\mu}_n\}$ is second order optimal for the sequence of performance measures $\{\hat{J}^{(n)}\}$.*

Proof. We simply put together the criterion of Lemma 3.3.2 with the properties established in Propositions 3.3.3 and C.3.5. ☐

Remark 3.3.1. Although it provides a reasonably large class of second order optimal sequences, Theorem 3.3.6 unfortunately does not give a criterion for second order optimality.

Example 3.3.7. A trivial example of a suitable sequence $\{\eta_n\}$ is obtained by setting $\eta_n \equiv \eta$ for some constant $\eta \in (0, 1)$.

3.4 The Pre-limit Sequence

3.4.1 The Infinite Buffer Setting

We return to the original problem announced in Subsection 3.1.5. Let us repeat the definition of the sequence of performance measures defined in (3.1.6). For all $n \in \mathbb{N}$, the mapping $J^{(n)} : \mathcal{L}^{(n)}(m) \to \mathbb{L}^0_+(\Omega, \mathcal{F}, \mathbb{P})$ is given by

$$J^{(n)}(\mu) = \int_0^T \mathbf{1}_{\{\frac{1}{n}Q_t^{(n)}(\mu)>K\}}\, dt \text{ for } \mu \in \mathcal{L}^{(n)}(m). \tag{3.4.1}$$

Our objective is to asymptotically minimize the values of $\{J^{(n)}\}$ along admissible sequences of service disciplines.

Definition 3.4.1. A sequence $\{\mu_n\}$ is called an *admissible sequence* if $\mu_n \in \mathcal{L}^{(n)}(m)$ for all $n \in \mathbb{N}$.

The following definition captures the notion of asymptotic optimality in this context.

Definition 3.4.2. An admissible sequence $\{\mu^*_n\}$ is called *asymptotically optimal* for the sequence of performance measures $\{J^{(n)}\}$, if

$$\liminf_{n\to\infty}[J^{(n)}(\mu_n) - J^{(n)}(\mu^*_n)] \geq 0,\ a.s.$$

for any other admissible sequence $\{\mu_n\}$.

We have the following criterion for asymptotic optimality.

Lemma 3.4.3. *Let $\{J^*_n\}$ be a sequence of random variables such that for all admissible sequences $\{\mu_n\}$,*

$$\liminf_{n\to\infty}[J^{(n)}(\mu_n) - J^*_n] \geq 0,\ a.s.$$

and let $\{\mu^*_n\}$ be an admissible sequence such that

$$\lim_{n\to\infty}[J^{(n)}(\mu^*_n) - J^*_n] = 0,\ a.s. \tag{3.4.2}$$

Then, $\{\mu^*_n\}$ is asymptotically optimal for $\{J^{(n)}\}$.

Proof. Let $\{J^*_n\}$ and $\{\mu^*_n\}$ satisfy the assumptions of the lemma. Then, for all admissible sequences $\{\mu_n\}$, we have

$$\begin{aligned}\liminf_{n\to\infty}[J^{(n)}(\mu_n) - J^{(n)}(\mu^*_n)] &\geq \liminf_{n\to\infty}[J^{(n)}(\mu_n) - J^*_n] + \liminf_{n\to\infty}[J^*_n - J^{(n)}(\mu^*_n)] \\ &\geq \lim_{n\to\infty}[J^*_n - J^{(n)}(\mu^*_n)] = 0,\end{aligned}$$

with all the inequalities above holding almost surely. □

The next lemma is dedicated to establishing the sequence of asymptotic lower bounds for the sequence of performance measures $\{J^{(n)}\}$. First, we introduce the sequence $\{\tau^{(n)}, n \in \mathbb{N}\}$ of (right-continuous) inverse processes of the normalized exogenous arrival processes $\frac{1}{n}N^+(n\mathcal{I}_t(\lambda))$, i.e., for every $l \in \mathbb{R}_+$, the action of the random mapping $\tau^{(n)}$ on l is defined as

$$\tau^{(n)}(l) = \inf\left\{t \geq 0 : \frac{1}{n}N^+(n\mathcal{I}_t(\lambda)) > l\right\} \wedge T. \tag{3.4.3}$$

Lemma 3.4.4. *Let $\{\mu_n\}$ be an admissible sequence. Consider the sequence of random variables $\{J_n^*\}$ defined by*

$$J_n^* = T - \tau_*^{(n)}, \tag{3.4.4}$$

where

$$\tau_*^{(n)} = \tau^{(n)}\left(K + \frac{1}{n}N^-(nm)\right)$$

with $\tau^{(n)}$ defined as in (3.4.3). Then, for all $n \in \mathbb{N}$,

$$J^{(n)} \geq J_n^*, \text{ a.s.}$$

Proof. Let the admissible sequence $\{\mu_n\}$ be given and let us temporarily fix an index n. It suffices to show that $\frac{1}{n}Q_t^{(n)} > K$ for every $t > \tau_*^{(n)}$. Since the Poisson process is nondecreasing and $\mu_n \in \mathcal{L}^{(n)}(m)$, we have that for every $t \in [0, T]$,

$$N^-(n\mathcal{I}_t(\mu_n)) \leq N^-(mn).$$

By the same reasoning, for every $t > \tau_*^{(n)}$,

$$\frac{1}{n}N^+(n\mathcal{I}_t(\lambda)) > \frac{1}{n}N^+(n\mathcal{I}_{\tau_*^{(n)}}(\lambda)) > K + \frac{1}{n}N^-(nm),$$

where the equality follows from the definition of the mapping $\tau^{(n)}$. Putting it all together, for all $t > \tau_*^{(n)}$, we see that

$$\frac{1}{n}Q_t^{(n)} \geq \frac{1}{n}X_t^{(n)} > K + \frac{1}{n}N^-(nm) - \frac{1}{n}N^-(nm) = K.$$

□

The next goal is to find an admissible sequence that is asymptotically optimal. It is enough to construct an admissible sequence $\{\mu_n^*\}$ which satisfies the criterion (3.4.2) with random variables $\{J_n^*\}$ provided in (3.4.4).

Assumption 3.4.5. Let $\{\gamma_n\}$ be a sequence in $(0, 1)$, that satisfies

(i) $\sum_{n=1}^{\infty} \frac{1}{n^2\gamma_n^4} < \infty$,

(ii) $\sum_{n=1}^{\infty} \frac{1}{n^2(1-\gamma_n)^4} < \infty$.

We now define the sequence $\{\tilde{\mu}_n\}$ of admissible service policies. For all n, we set

$$\tilde{\mu}_n(t) = \begin{cases} 0 & \text{for } t \leq \tau(\gamma_n K), \\ \lambda_t & \text{for } \tau(\gamma_n K) < t \leq \tau(\gamma_n K + m), \\ 0 & \text{for } \tau(\gamma_n K + m) < t \leq T. \end{cases} \tag{3.4.5}$$

This sequence of service policies generates the following sequence of netput processes

$$X_t^{(n)}(\tilde{\mu}_n) = \begin{cases} N^+(n\mathcal{I}_t(\lambda)) & t \in [0, \tau(\gamma_n K)], \\ N^+(n\mathcal{I}_t(\lambda)) - N^-(n(\mathcal{I}_t(\lambda) - \gamma_n K)) & t \in (\tau(\gamma_n K), \tau(\gamma_n K + m)], \\ N^+(n\mathcal{I}_t(\lambda)) - N^-(nm) & t \in (\tau(\gamma_n K + m), T]. \end{cases}$$

The corresponding queue length processes are given by $Q^{(n)}(\tilde{\mu}_n) = \Gamma(X^{(n)}(\tilde{\mu}_n))$. Our objective is to verify condition (3.4.2) for the admissible sequence $\{\tilde{\mu}_n\}$.

Further analysis is simplified by Assumption 3.4.5(i). Its role is to cause an almost sure absence of regulation in the limit, as can be seen in the following proposition which follows from Corollary C.4.4.

Proposition 3.4.6. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5(i). Then, $Q_t^{(n)}(\tilde{\mu}_n) = X_t^{(n)}(\tilde{\mu}_n)$ for all $t \in [\tau(\gamma_n K), \tau(\gamma_n K + m)]$, for all but finitely many n, almost surely.*

We have resolved that Assumption 3.4.5(i) is indeed a sufficient condition for eventual absence of reflection on the segment $[\tau(\gamma_n K), \tau(\gamma_n K+m)]$. Next, we verify that the complete Assumption 3.4.5 is a sufficient condition for the sequence $\{\gamma_n\}$ to induce an asymptotically optimal $\{\tilde{\mu}_n\}$.

Proposition 3.4.7. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5. Then*

$$\mathbb{P}\left[\sup_t \frac{1}{n} Q_t^{(n)}(\tilde{\mu}_n) \leq K, ev.\right] = 1,$$

with the supremum in the previous display taken over all $t \in [\tau(\gamma_n K), \tau(\gamma_n K + m)]$.

Proof. Proposition C.4.7 implies

$$\sum_{n=1}^{\infty} \mathbb{P}\left[\inf_t \frac{1}{n} Q_t^{(n)}(\tilde{\mu}_n) > K\right] \leq C \left[\sum_{n=1}^{\infty} \frac{1}{n^2(1-\gamma_n)^4} + \sum_{n=1}^{\infty} \frac{1}{n^2\gamma_n^4}\right] < \infty,$$

and the announced claim is a simple consequence of the Borel-Cantelli Lemma. ☐

Remark 3.4.1. It is of obvious interest to show that the assumptions imposed on $\{\gamma_n\}$ are not void. To produce a trivial example, we can set $\gamma_n = \frac{1}{2}$ for all n.

For an example of $\{\gamma_n\}$ converging to 1, let $\gamma_n = 1 - \frac{1}{2}n^{-1/5}$ for all $n \in \mathbb{N}$. Then

$$\begin{aligned}\sum_{n=1}^{\infty} \frac{1}{n^2\gamma_n^4} &= 16 \sum_{n=1}^{\infty} \frac{1}{n^2 n^{-4/5}(2n^{1/5}-1)^4} \\ &= 16 \sum_{n=1}^{\infty} \frac{1}{n^{6/5}(2n^{1/5}-1)^4} \leq 16 \sum_{n=1}^{\infty} \frac{1}{n^{6/5}} < \infty\end{aligned}$$

and

$$\sum_{n=1}^{\infty} \frac{1}{n^2(1-\gamma_n)^4} = 16 \sum_{n=1}^{\infty} \frac{1}{n^2 n^{-4/5}} = 16 \sum_{n=1}^{\infty} \frac{1}{n^{6/5}} < \infty.$$

Hence, Assumption 3.4.5 is satisfied by the proposed sequence.

Theorem 3.4.8. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5. Then we have almost surely, for almost every n,*

$$J^{(n)}(\tilde{\mu}_n) = T - \tau_*^{(n)}.$$

In other words, the sequence of service policies $\{\tilde{\mu}_n\}$ is asymptotically optimal for $\{J^{(n)}\}$.

Proof. We use Lemma C.4.1 to rule out any contribution to the cost before time $\tau(\gamma_n K)$ and Proposition 3.4.7 to eliminate cost aggregated over the interval $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$, both for sufficiently large n. The result is the equality

$$J^{(n)}(\tilde{\mu}_n) = \int_{\tau(\gamma_n K+m)}^{T} \mathbf{1}_{\{\frac{1}{n}Q_t^{(n)}(\tilde{\mu}_n)>K\}}\, dt, \textit{ a.s.}, \tag{3.4.6}$$

which holds for sufficiently large n. In view of Corollary C.4.4, we rewrite (3.4.6) as

$$\begin{aligned}J^{(n)}(\tilde{\mu}_n) &= \int_{\tau(\gamma_n K+m)}^{T} \mathbf{1}_{\{\frac{1}{n}(N^+(n\mathcal{I}_t(\lambda))-N^-(nm))>K\}}\, dt \\ &= \left(T - \tau(\gamma_n K + m) \vee \tau_*^{(n)}\right)^+ \leq \left(T - \tau_*^{(n)}\right)^+.\end{aligned}$$

We recognize the right-hand side as the lower bound for the mapping $J^{(n)}$ obtained in Lemma 3.4.4, which completes the proof. □

In general, we cannot make any claims on the almost sure limit of the sequence of minimal penalties from the previous theorem. However, under slight additional assumptions stated below, the sequence $\{J^{(n)}(\tilde{\mu}_n)\}$ does have a limit, and that limit equals the optimal penalty in the fluid case J^*, defined in (3.2.8). The appropriate additional assumption is that the cumulative arrival rate $\mathcal{I}(\lambda)$ is strictly increasing in the neighborhood of $\tau(K+m)$. The precise claims and their proofs are exhibited Appendix C.4. Appendix C.5 contains a comparison between the optimal classes we obtained in the previous three sections.

3.4.2 Consequences in the Finite Buffer Setting

In this subsection, $Q^{(n)}, L^{(n)}$ and $U^{(n)}$ are defined as in Subsections 3.1.2 and 4.3.1

As in the discussion of the fluid regime, our strategy consists of two steps - finding an (asymptotic) lower bound on the performance measures and then proving that a particular class of controls attains that lower bound.

Lemma 3.4.9. *For every n, and for every $\mu \in \mathcal{L}^{(n)}(m)$, we have*

$$J_F^{(n)}(\mu) \geq \left[\frac{1}{n}N^+(n\mathcal{I}_T(\lambda)) - \frac{1}{n}N^-(nm) - K\right]^+, \text{ a.s.},$$

where $J_F^{(n)} : \mathcal{L}^{(n)}(m) \to \mathbb{L}^0_+(\Omega, \mathcal{F}, \mathbb{P})$ is defined in (3.1.4).

Proof. Let us temporarily fix an index n and a service discipline $\mu \in \mathcal{L}^{(n)}(m)$. Using expression (3.1.3) and the facts that $L^{(n)}(\mu) \geq 0$ and $Q^{(n)}(\mu) \leq nK$, we get

$$J_F^{(n)}(\mu) = \frac{1}{n}U_T^{(n)}(\mu) \geq \frac{1}{n}N^+(n\mathcal{I}_T(\lambda)) - \frac{1}{n}N^-(n\mathcal{I}_T(\mu)) - K. \tag{3.4.7}$$

Since $\mu \in \mathcal{L}^{(n)}(m)$, we have that $\mathcal{I}_T(\mu) \leq m, a.s.$ This bound, combined with the monotonicity of the Poisson process N^- and inequality (3.4.7), implies that

$$J_F^{(n)}(\mu) = \frac{1}{n}U_T^{(n)}(\mu) \geq \frac{1}{n}N^+(n\mathcal{I}_T(\lambda)) - \frac{1}{n}N^-(nm) - K.$$

The fact that $U_T^{(n)} \geq 0$ by definition wraps up the proof. □

The following theorem is the main result of this section.

Theorem 3.4.10. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5 and let $\{\tilde{\mu}_n\}$ be the corresponding sequence of polices, as defined in (3.4.5). Then we have almost surely, for all but finitely many n*

$$J_F^{(n)}(\tilde{\mu}_n) = \left[\frac{1}{n}N^+(n\mathcal{I}_T(\lambda)) - \frac{1}{n}N^-(nm) - K\right]^+,$$

i.e., the sequence of service policies $\{\tilde{\mu}_n\}$ is asymptotically optimal for $\{J_F^{(n)}\}$.

Proof. We start by expressing for every n the length of the queue, using Theorem 1.4 of [KLRS06]. For every $t \in [0, T]$, we have

$$Q_t^{(n)}(\tilde{\mu}_n) = \Gamma(X^{(n)}(\tilde{\mu}_n))_t - \sup_{0 \leq s \leq t}\left[\left(\Gamma(X^{(n)}(\tilde{\mu}_n))_s - K\right)^+ \wedge \inf_{s \leq u \leq t}\Gamma(X^{(n)}(\tilde{\mu}_n))_u\right]. \tag{3.4.8}$$

Lemma C.4.1 and Propositions C.4.3 and 3.4.7 allow us to conclude that

$$\Gamma(X^{(n)}(\tilde{\mu}_n))_t \in [0, K], \text{ for all } t \leq \tau_*^{(n)},$$

almost surely, for all but finitely many n. Hence, evaluated at any instant t preceding $\tau_*^{(n)}$ the equation (3.4.8) has the form

$$Q_t^{(n)}(\tilde{\mu}_n) = \Gamma(X^{(n)}(\tilde{\mu}_n))_t - \sup_{0\le s\le t}\left[0 \wedge \inf_{s\le u\le t}\Gamma(X^{(n)}(\tilde{\mu}_n))_u\right] = \Gamma(X^{(n)}(\tilde{\mu}_n))_t. \tag{3.4.9}$$

We conclude that $L_{\tau_*^{(n)}}^{(n)}(\tilde{\mu}_n) = U_{\tau_*^{(n)}}^{(n)}(\tilde{\mu}_n) = 0$, almost surely, for all but finitely many n. Moreover, since there is no remaining service on $(\tau_*^{(n)}, T]$, it must be that $L_T^{(n)}(\tilde{\mu}_n) = 0$, almost surely, for all but finitely many n.

Returning again to (3.4.8) we express the total amount of regulation on $[0,T]$ as

$$U_T^{(n)}(\tilde{\mu}_n) = \sup_{\tau_*^{(n)}\le s\le T}\left[(\Gamma(X^{(n)}(\tilde{\mu}_n))_s - K)^+ \wedge \inf_{s\le u\le t}\Gamma(X^{(n)}(\tilde{\mu}_n))_u\right].$$

Using Proposition C.4.3 this translates into

$$U_T^{(n)}(\tilde{\mu}_n) = \sup_{\tau_*^{(n)}\le s\le T}\left[(X_s^{(n)}(\tilde{\mu}_n) - K)^+ \wedge \inf_{s\le u\le t} X_u^{(n)}(\tilde{\mu}_n)\right].$$

As there is no service in the region across which the supremum is taken in the expression above, the netput process is nondecreasing and we get

$$U_T^{(n)}(\tilde{\mu}_n) = \sup_{\tau_*^{(n)}\le s\le T}\left[(X_s^{(n)}(\tilde{\mu}_n) - K)^+ \wedge X_s^{(n)}(\tilde{\mu}_n)\right] = \sup_{\tau_*^{(n)}\le s\le T}(X_s^{(n)}(\tilde{\mu}_n) - K)^+.$$

Again, since the netput process is nondecreasing, we obtain

$$U_T^{(n)}(\tilde{\mu}_n) = (X_T^{(n)}(\tilde{\mu}_n) - K)^+ = (N^+(n\mathcal{I}_T(\lambda)) - N^-(nm) - K)^+. \tag{3.4.10}$$

Dividing the equality in (3.4.10) throughout by n completes the proof. □

Chapter 4

The Tandem Network

4.1 The Control Problem - Finite Buffers

We next focus on the two-station tandem queueing network, with each station having a single server and a finite buffer capacity. The modeling assumption is that the process of exogenous arrivals into the first station, as well as the processes of potential departures are nonhomogeneous Poisson. We implement these assumptions as follows:

Let N_1^+, N_1^- and N_2^- be independent unit Poisson processes. Moreover, the time-varying rates are

- $\lambda \in \mathbb{L}_+^1[0,T]$ - the known, deterministic exogenous arrival rate;
- μ - the service rate at the first station, which we are at liberty to choose among all non-anticipating μ such that $\mathcal{I}_T(\mu) \leq m$. Appendix B.3 contains the formal definition of the space of all non-anticipating service disciplines. In the present section, we refer to the set of all such service disciplines satisfying the constraint $\mathcal{I}_T(\mu) \leq m$ as *admissible.* Obviously, all deterministic service disciplines μ satisfying $\mathcal{I}_T(\mu) \leq m$ are admissible.
- $\mu^2 \in \mathbb{L}_+^1[0,T]$ - the service rate at the second station, assumed to be deterministic and known.

The capacities of the two finite buffers are denoted by $K_1, K_2 \in \mathbb{R}_+ \cup \{\infty\}$. We permit the buffer capacities to assume the value of ∞ in order to allow the extremal infinite buffer case for any or both of the stations.

Next, we formally describe the dynamics of the system. As all the other values are given, we will denote all the derived processes as functions depending on the admissible control μ. The netput process in the first station per the above assumptions is

$$X^{(1)}(\mu) = N_1^+(\mathcal{I}(\lambda)) - N_1^-(\mathcal{I}(\mu)).$$

The defining equations of the two-sided reflection mapping Γ^{K_1} (see Definition 1.2.2) applied to the process $X^{(1)}(\mu)$ yield the queue length

$$Q^{(1)}(\mu) = X^{(1)}(\mu) + L^{(1)}(\mu) - U^{(1)}(\mu), \tag{4.1.1}$$

where $L^{(1)}(\mu)$ and $U^{(1)}(\mu)$ are the regulator maps associated with $X^{(1)}(\mu)$ and K_1 in the sense of Definition 1.2.2.

The departure process from the first station is then

$$N_1^+(\mathcal{I}(\lambda)) - Q^{(1)}(\mu) - U^{(1)}(\mu) = N_1^-(\mathcal{I}(\mu)) - L^{(1)}(\mu). \tag{4.1.2}$$

As the process in (4.1.2) is also the arrival process into the second station, the netput process for the second station reads as

$$X^{(2)}(\mu) = N_1^-(\mathcal{I}(\mu)) - L^{(1)}(\mu) - N_2^-(\mathcal{I}(\mu^2)),$$

and the queue length is

$$Q^{(2)}(\mu) = X^{(2)}(\mu) + L^{(2)}(\mu) - U^{(2)}(\mu), \tag{4.1.3}$$

with the regulator processes $L^{(2)}(\mu)$ and $U^{(2)}(\mu)$ associated with $X^{(2)}(\mu)$ and K_2 as in Definition 1.2.2.

Due to the finiteness of the two buffers in the system, a certain number of customers may be lost - potentially causing a cost to the manager of the system. The number of customers lost during the interval $[0, T]$, based on expressions (4.1.1) and (4.1.3), is

$$U_T^{(1)}(\mu) + U_T^{(2)}(\mu)$$

for any admissible service discipline μ. We want to find a way to minimize this cost by varying the service discipline μ.

4.2 The Uniform Acceleration Scheme

As is expected, we do not look at the original control problem *per se*. Instead, we consider a sequence of systems with the same network topology as the original system, but with uniformly accelerated rates. We use the same symbols for the buffer sizes and the basic rates of Poisson processes appearing in the sequence of systems as we did in the original model encapsulated in (4.1.1) and (4.1.3).

To any fixed index n, there corresponds a different space $\mathcal{L}^{(n)}(m)$ of admissible service disciplines. The construction of these spaces is described in Appendix B.3. A precise description of the model for the sequence of tandems is next.

For every n, we list relevant stochastic processes expressed as functions of admissible service disciplines $\mu \in \mathcal{L}^{(n)}(m)$. The dynamics in the first station is described by the netput process

$$X^{(1,n)}(\mu) = N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu)),$$

which, upon the application of the two-sided reflection map constraining the queue to the strip $[0, nK_1]$ in accordance with the scaled buffer capacity nK_1, gives rise to the queue length process

$$Q^{(1,n)}(\mu) = X^{(1,n)}(\mu) + L^{(1,n)}(\mu) - U^{(1,n)}(\mu), \tag{4.2.1}$$

with $L^{(1,n)}(\mu)$ and $U^{(1,n)}(\mu)$ being the regulator maps associated with $X^{(1,n)}(\mu)$ and nK_1 in the sense of Definition 1.2.2.

The netput process in the second station is

$$X^{(2,n)}(\mu) = N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu) - N_2^-(n\mathcal{I}(\mu^2)),$$

and the queue length is

$$Q^{(2,n)}(\mu) = X^{(2,n)}(\mu) + L^{(2,n)}(\mu) - U^{(2,n)}(\mu). \tag{4.2.2}$$

The regulator maps $L^{(2,n)}(\mu)$ and $U^{(2,n)}(\mu)$ are chosen in the sense of Definition 1.2.2 to correspond to the two-sided reflection map on $[0, nK_2]$ applied to the process $X^{(2,n)}(\mu)$.

For every n, the mapping $J_F^{(n)} : \mathcal{L}^{(n)}(m) \to \mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$, defined by

$$J_F^{(n)}(\mu) = \frac{1}{n}\left(U_T^{(1,n)}(\mu) + U_T^{(2,n)}(\mu)\right) \text{ for every } \mu, \tag{4.2.3}$$

measures the performance of the system in terms of the normalized number of lost jobs in both stations during $[0, T]$.

4.3 The Auxiliary Problem - Infinite Buffers

We move on to the tandem system in which the buffers in both stations are infinite, so that the only restriction imposed on the queue length processes is the nonnegativity requirement. The queue length processes are

$$\begin{aligned} Q^{(1,n)}(\mu) &= N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu)) + L^{(1,n)}(\mu), \\ Q^{(2,n)}(\mu) &= N_1^-(n\mathcal{I}(\mu)) - L^{(1,n)}(\mu) - N_2^-(n\mathcal{I}(\mu^2)) + L^{(2,n)}(\mu), \end{aligned} \tag{4.3.1}$$

where the processes $L^{(1,n)}(\mu)$ and $L^{(2,n)}(\mu)$ are the regulating terms arising in Definition 1.2.1 of the one-dimensional reflection map. Comparing (4.3.1) to (4.2.1) and (4.1.3), we see that in the extreme case of $K_1 = K_2 = \infty$, the expressions for the queue lengths coincide, as is to be expected.

Similarly to the notation in the single station case, $\{\tau^{(1,n)}, n \in \mathbb{N}\}$ represent the (right-continuous) inverse processes of the normalized exogenous arrival processes $\frac{1}{n}N_1^+(n\mathcal{I}_t(\lambda))$, i.e.,

$$\tau^{(1,n)}(l) = \inf\left\{t \geq 0 : \frac{1}{n}N_1^+(n\mathcal{I}_t(\lambda)) > l\right\} \wedge T, \text{ for all } l \in \mathbb{R}_+.$$

4.3.1 The Performance Measure

Let the two positive constants K_1 and K_2 now denote the thresholds in the two queues in the system. For every $n \in \mathbb{N}$, we define the performance measure $J^{(n)} : \mathcal{L}^{(n)}(m) \longrightarrow \mathbb{L}_+^0(\Omega, \mathcal{F}, \mathbb{P})$ by

$$J^{(n)}(\mu) = \int_0^T \left[\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,n)}(\mu)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,n)}(\mu)>K_2\}}\right] dt, \text{ for } \mu \in \mathcal{L}^{(n)}(m). \tag{4.3.2}$$

Occasionally it will be more convenient to use the following equivalent formulation of $J^{(n)}$

$$J^{(n)}(\mu) = meas\left\{\frac{1}{n}Q_t^{(1,n)}(\mu) > K_1\right\} + meas\left\{\frac{1}{n}Q_t^{(2,n)}(\mu) > K_2\right\} \tag{4.3.3}$$

for $\mu \in \mathcal{L}^{(n)}(m)$.

We are interested in minimizing the values of the above performance measures across all sequences of admissible controls in a suitable (asymptotic) manner to be described later on (see Definition 4.6.1).

Remark 4.3.1. We choose to think about the cost structure of (4.3.2) in terms of a unit penalty that is accumulated over the period $[0, T]$ whenever a queue exceeds the buffer level. If both queues are above their thresholds simultaneously, then the penalties are added up. From now on we will occasionally refer to the integrand in (4.3.2) or - depending on the context - the performance measure itself as the penalty.

Remark 4.3.2. The similarities between this auxiliary infinite-buffer problem and the main control problem proposed in the previous section are not hard to see. A natural principle that we expect to arise in solving both problems is to try to restrict the queue lengths so that whenever possible

(a) they are below or at the threshold level in the infinite buffer case;

(b) the downward pushing is avoided in the finite buffer case.

However, it turns out that the controls that perform well in one case are not guaranteed to perform well in the other setting - the situation is more complicated than that.

We intend to solve the control problem for the infinite buffer case - first in the context of the fluid-limit problem in the following section. We then proceed to solve the fluid-limit problem associated with the finite-buffer case through simple analogy with the infinite-buffer case. This approach does not extend to our central problem of asymptotic optimality covered in Sections 4.6 and 4.8, but the fluid-limit optimal classes are already informative enough to facilitate the solutions of both the finite-buffer and the infinite-buffer asymptotic control problems.

4.4 Fluid Limit Analysis in the Infinite-Buffer Case

Given a fixed deterministic service discipline μ, by (4.3.1), Theorem 1.2.5 and the continuity of the one-sided regulator map in the uniform topology (see Theorem 13.5.1 in [Whi02b]), we have

$$\begin{aligned} &\frac{1}{n}Q^{(1,n)}(\mu) \to \bar{Q}^{(1)}(\mu) := \mathcal{I}(\lambda) - \mathcal{I}(\mu) + \bar{L}^{(1)}(\mu), \\ &\frac{1}{n}Q^{(2,n)}(\mu) \to \bar{Q}^{(2)}(\mu) := \mathcal{I}(\mu) - \bar{L}^{(1)}(\mu) - \mathcal{I}(\mu^2) + \bar{L}^{(2)}(\mu), \end{aligned} \tag{4.4.1}$$

almost surely, in the uniform topology, where

$$\begin{aligned} \bar{L}^{(1)}(\mu) &= \sup_{s\leq \cdot}[-\mathcal{I}_s(\lambda - \mu)], \\ \bar{L}^{(2)}(\mu) &= \sup_{s\leq \cdot}[-\mathcal{I}_s(\mu) + \bar{L}^{(1)}_s(\mu) + \mathcal{I}_s(\mu^2)] \end{aligned}$$

are the regulator maps corresponding to $\mathcal{I}(\lambda) - \mathcal{I}(\mu)$ and $\mathcal{I}(\mu) - \mathcal{I}(\mu^2) - \bar{L}^{(1)}(\mu)$, respectively, for the one-sided reflection map Γ of Definition 1.2.1.

4.4.1 Performance Measure

In order to define the performance measure analogous to the ones introduced in Subsection 4.3.1, we first need to introduce the set of admissible service disciplines. The service disciplines we allow in the context of the fluid limit are deterministic and their integrals must not exceed the given bound m. Again, we denote the set of admissible service disciplines by

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

The fluid-limit variant of the performance measures from (4.3.2) is defined as $\bar{J} : \bar{\mathcal{L}}(m) \to \mathbb{R}_+$, with

$$\bar{J}(\mu) = \text{meas}\{t \in [0,T] : \bar{Q}^{(1)}_t(\mu) > K_1\} + \text{meas}\{t \in [0,T] : \bar{Q}^{(2)}_t(\mu) > K_2\}, \text{ for every } \mu \in \bar{\mathcal{L}}(m).$$

Our goal is to identify service disciplines that minimize $\bar{J}$, provided that such service disciplines exist. We then have the larger goals of relating the fluid limit problem with the performance measures $J^{(n)}$ of (4.3.2) to the finite-buffer problem. We will first determine a lower bound on $\bar{J}$, and follow with a criterion identifying the service disciplines achieving that lower bound.

4.4.2 The Lower Bound

A convenient yardstick for the performance of the system just described is the performance of the corresponding "pooled" queue, which is obtained by omitting the first station in the tandem system. Since the rate of exogenous arrivals λ and the rate of service in the second station μ^2 are given, the queue length $\bar{Q}^P$ in this "pooled" system reads as

$$\bar{Q}^P = \bar{X}^P + \bar{L}^P, \tag{4.4.2}$$

with $\bar{X}^P = \mathcal{I}(\lambda - \mu^2)$ and $\bar{L}^P = \sup_{s\leq \cdot}[-\mathcal{I}_s(\lambda - \mu^2)]$. Note that $\bar{Q}^P$ does not depend on the service discipline at the first station.

Lemma 4.4.1. *For every $\mu \in \bar{\mathcal{L}}(m)$ and every $t \in [0,T]$, we have that $\bar{L}^{(2)}_t(\mu) \geq \bar{L}^P_t$ and*

$$\bar{Q}^{(1)}_t(\mu) + \bar{Q}^{(2)}_t(\mu) \geq \bar{Q}^P_t. \tag{4.4.3}$$

Proof. Using (4.4.1) for every μ and t, the left-hand side of (4.4.3) can be rewritten as

$$\begin{aligned}\bar{Q}^{(1)}_t(\mu) + \bar{Q}^{(2)}_t(\mu) &= \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu) + \bar{L}^{(1)}_t(\mu) + \mathcal{I}_t(\mu) - \bar{L}^{(1)}_t(\mu) - \mathcal{I}_t(\mu^2) + \bar{L}^{(2)}_t(\mu)\\ &= \mathcal{I}_t(\lambda - \mu^2) + \bar{L}^{(2)}_t(\mu)\\ &= \bar{X}^P_t + \bar{L}^{(2)}_t(\mu).\end{aligned}$$

Therefore, it suffices to prove $\bar{L}^{(2)}_t(\mu) \geq \bar{L}^P_t$. By definition,

$$\bar{L}^{(2)}_t(\mu) = \sup_{s\leq t}[-\mathcal{I}_s(\mu) + \bar{L}^{(1)}_s(\mu) + \mathcal{I}_s(\mu^2)] = \sup_{s\leq t}[\bar{Q}^{(1)}_s - \mathcal{I}_s(\lambda) + \mathcal{I}_s(\mu^2)].$$

Since the queue length $\bar{Q}^{(1)}$ is nonnegative,

$$\bar{L}^{(2)}_t(\mu) \geq \sup_{s\leq t}[-\mathcal{I}_s(\lambda) + \mathcal{I}_s(\mu^2)] = \bar{L}^P_t. \tag{4.4.4}$$

□

Now that we know that the length of the pooled queue is a lower bound for the sum of the lengths of the two queues in tandem, we can get a lower bound on the aggregated penalty up to any time $t \in [0,T]$. Further exposition will be facilitated by introducing the mappings $\bar{J} : [0,T] \times \bar{\mathcal{L}}(m) \to \mathbb{R}_+$ and $\bar{J}^P : [0,T] \to \mathbb{R}_+$ as follows

$$\begin{aligned}\bar{J}(t,\mu) &= \text{meas}\,\{s \in [0,t] : \bar{Q}^{(1)}_s(\mu) > K_1\} + \text{meas}\,\{s \in [0,t] : \bar{Q}^{(2)}_s(\mu) > K_2\},\\ \bar{J}^P(t) &= \text{meas}\,\{s \in [0,t] : \bar{Q}^P_s > K_1 + K_2\}.\end{aligned} \tag{4.4.5}$$

In particular, we have that $\bar{J}(T,\mu) = \bar{J}(\mu)$.

Lemma 4.4.2. *For every admissible service discipline μ and every $t \in [0,T]$, we have*

$$\bar{J}(t,\mu) \geq \bar{J}^P(t).$$

Proof. Let us temporarily fix a service discipline μ and a time t. For all $s \leq t$ such that $\bar{Q}_s^P > K_1 + K_2$, by Lemma 4.4.1 we have that $\bar{Q}_s^{(1)}(\mu) + \bar{Q}_s^{(2)}(\mu) > K_1 + K_2$. Hence, at least one of the inequalities $\bar{Q}_s^{(1)}(\mu) > K_1$ and $\bar{Q}_s^{(2)}(\mu) > K_2$ must hold true, which completes the proof. □

The previous lemma directs our further investigation as follows. Let us imagine a myopic observer who does not see the tandem system in its entirety - he/she only sees the arrival stream into the first station and the departures from the second station. We aim at finding these service disciplines in the first station that have the property that the said observer would not be able to tell whether he/she is looking at the pooled queue or at a tandem system with a hidden first station. This is feasible only "locally", i.e., until the constraint on the amount of service available is reached, which prevents communication between the two stations, and hence, further emulation of the pooled queue by the tandem system. If this is achieved, then the lower bound generated by the pooled queue from the last lemma is matched and we have an optimal control. Thus, we need a different lower bound on the penalty accumulated from the time when even if all the available service is utilized, the first queue must cross over the threshold. Let us recall that $\tau : \mathbb{R}_+ \to [0, T]$ is the (right-continuous) inverse mapping of $\mathcal{I}(\lambda) : [0, T] \to \mathbb{R}_+$ as defined in (3.2.7).

Lemma 4.4.3. *For all $\mu \in \bar{\mathcal{L}}(m)$,*

$$\bar{J}(T, \mu) - \bar{J}(\tau(K_1 + m), \mu) = \bar{J}(\mu) - \bar{J}(\tau(K_1 + m), \mu) \geq T - \tau(K_1 + m).$$

Proof. By the definition of τ, for every $\mu \in \bar{\mathcal{L}}(m)$ we have

$$\bar{Q}_t^{(1)}(\mu) \geq \mathcal{I}_t(\lambda - \mu) \geq \mathcal{I}_t(\lambda) - m > K_1, \text{ for every } t > \tau(K_1 + m). \tag{4.4.6}$$

From the definition in (4.4.5),

$$\begin{aligned} \bar{J}(\mu) - \bar{J}(\tau(K_1 + m), \mu) &\geq \text{meas}\,\{t \in [\tau(K_1 + m), T] : \bar{Q}_t^{(1)}(\mu) > K_1\} \\ &= T - \tau(K_1 + m). \end{aligned} \tag{4.4.7}$$

□

Combining Lemmas 4.4.2 and 4.4.3, we get the following lower bound on the performance measure $\bar{J}$.

Corollary 4.4.4. *For all admissible service disciplines μ, we have*

$$\bar{J}(\mu) \geq \bar{J}^P(\tau(K_1 + m)) + T - \tau(K_1 + m).$$

4.4.3 A Class of Fluid-Optimal Disciplines

For the sake of completeness, let us state the formal definition of a fluid-optimal discipline.

Definition 4.4.5. We say that a service discipline $\mu^* \in \bar{\mathcal{L}}(m)$ is *fluid optimal* for the performance measure $\bar{J}$ if $\bar{J}(\mu) \geq \bar{J}(\mu^*)$ for all $\mu \in \bar{\mathcal{L}}(m)$.

Let us denote by $\bar{\mathcal{L}}^*(m)$ the subset of $\bar{\mathcal{L}}(m)$ containing all service disciplines μ satisfying the following conditions:

$\bar{\mathcal{L}}^*(m)$: 1. $\bar{L}_T^{(1)}(\mu) = 0$ or, equivalently, $\mathcal{I}_t(\lambda) \geq \mathcal{I}_t(\mu)$, for all t.

$\bar{\mathcal{L}}^*(m)$: 2. For all $t \leq \tau(K_1 + m)$,

$$\bar{Q}_t^P \leq K_1 + K_2 \Rightarrow \mathcal{I}_t(\lambda) - K_1 \leq \mathcal{I}_t(\mu) \leq K_2 + \mathcal{I}_t(\mu^2) - \bar{L}_t^P. \tag{4.4.8}$$

$\bar{\mathcal{L}}^*(m)$: 3. For all t, $\mathcal{I}_t(\lambda) - K_1 \leq \mathcal{I}_t(\mu)$ or $\mathcal{I}_t(\mu) \leq K_2 + \mathcal{I}_t(\mu^2) - \bar{L}_t^P$.

In words, the assumptions above mean that ($\bar{\mathcal{L}}_m^*$: 1) there is never reflection off zero in the first queue, ($\bar{\mathcal{L}}_m^*$: 2) whenever the pooled queue is below its threshold $K_1 + K_2$, both queues in the tandem must also be below their thresholds and ($\bar{\mathcal{L}}_m^*$: 3) at no time can both queues in the tandem be above their thresholds. We claim the following.

Proposition 4.4.6. *All $\mu \in \bar{\mathcal{L}}^*(m)$ are fluid-optimal.*

Proof. Let μ be a service discipline in $\bar{\mathcal{L}}^*(m)$. By condition ($\bar{\mathcal{L}}^*(m)$: 3), for every t we have

$$\bar{J}(t, \mu) = \text{meas}\,\{s \in [0, t] : \bar{Q}_s^{(1)}(\mu) > K_1 \text{ or } \bar{Q}_s^{(2)}(\mu) > K_2\}.$$

Using the remaining defining conditions on $\bar{\mathcal{L}}^*(m)$, the penalty accumulated up to time $\tau(K_1+m)$ can be further rewritten as

$$\begin{aligned} \bar{J}(\tau(K_1 + m), \mu) &= \text{meas}\,\{s \in [0, \tau(K_1 + m)] : \bar{Q}_s^{(1)}(\mu) > K_1 \text{ or } \bar{Q}_s^{(2)}(\mu) > K_2\} \\ &= \bar{J}^P(\tau(K_1 + m)), \end{aligned}$$

where the first equality follows from ($\bar{\mathcal{L}}_m^*$: 3) and the second equality results from ($\bar{\mathcal{L}}_m^*$: 2). On the other hand, for every $t > \tau(K_1 + m)$, we necessarily have $\bar{Q}_t^{(1)}(\mu) > K_1$ (see the proof of Lemma 4.4.3). By condition ($\bar{\mathcal{L}}^*(m)$: 3), hence, $\bar{Q}_t^{(2)}(\mu) \leq K_2$. Altogether, we conclude that

$$\bar{J}(\mu) = \bar{J}^P(\tau(K_1 + m)) + T - \tau(K_1 + m),$$

i.e., the lower bound established in Corollary 4.4.4 is met. □

It is of obvious interest to deliver a particular fluid-optimal service discipline. We dedicate the following subsection to that task.

4.4.4 A fluid-optimal discipline

The strategy we wish to explore is to serve at full speed, while avoiding serving in vain in the first station, as well as incurring penalty in the second station. It is suitable to use the length of the pooled queue to implement this strategy.

Let us begin by introducing $\mu^{aux} \in \mathbb{L}^1_+[0,T]$ as

$$\mu^{aux}_t = \begin{cases} \lambda_t & \text{if } \bar{Q}^P_t < K_2 \\ \mu^2_t & \text{if } \bar{Q}^P_t \geq K_2. \end{cases} \tag{4.4.9}$$

In order to define an admissible discipline, we need to enforce the cap m on the cumulative service available. So let

$$t^* = \inf\{t > 0 : \mathcal{I}_t(\mu^{aux}) \geq m\} \wedge T. \tag{4.4.10}$$

We claim that the admissible service discipline μ^* given by

$$\mu^* = \mu^{aux}\mathbf{1}_{[0,t^*]} \tag{4.4.11}$$

is fluid optimal. The left graph in Figure 4.1 shows the evolution of the pooled queue length, while the right graph shows the state of the tandem system in the $(\bar{Q}^{(1)}, \bar{Q}^{(2)})$ quarter-plane. The intervals of time the pooled queue spends above its threshold $K = K_1 + K_2$ correspond to the intervals of time the state process spends in the region $(K_1, \infty) \times \{K_2\}$. This correspondence between the pooled queue and the tandem is valid until the total service rendered in the first station reaches its constraint. The point at which the constraint is reached is exactly the point on the right-hand graph at which the state process starts moving in the "south-east" direction.

The netput process of the first queue in tandem at times t preceding t^* can be represented as

$$\bar{X}^{(1)}_t(\mu^*) = \mathcal{I}_t(\lambda - \mu^*) = \int_0^t (\lambda_s - \mu^2_s)\mathbf{1}_{\{\bar{Q}^P_s = K_2\}}\, ds + \int_0^t (\lambda_s - \mu^2_s)\mathbf{1}_{\{\bar{Q}^P_s > K_2\}}\, ds. \tag{4.4.12}$$

Thanks to the fact that the length of the pooled queue has the following integral form

$$\bar{Q}^P_t = \int_0^t (\lambda_s - \mu^2_s)\mathbf{1}_{\{\bar{Q}^P_s > 0\}}\, ds + \int_0^t (\lambda_s - \mu^2_s)^+\mathbf{1}_{\{\bar{Q}^P_s = 0\}}\, ds, \tag{4.4.13}$$

we can, setting

$$f_s = (\lambda_s - \mu^2_s)\mathbf{1}_{\{\bar{Q}^P_s > 0\}} + (\lambda_s - \mu^2_s)^+\mathbf{1}_{\{\bar{Q}^P_s = 0\}} \qquad \text{and} \qquad F = \bar{Q}^P,$$

employ results (D.2.4) and (D.2.5) in Appendix D.2 and conclude that the first integral in (4.4.12) disappears, while the second one reduces to $(\bar{Q}^P_t - K_2)^+$. This is a nonnegative value,

Figure 4.1: The Tandem System: A fluid-optimal policy in the case of infinite buffers with $K_1 = 0.2$ and $K_2 = 0.1$

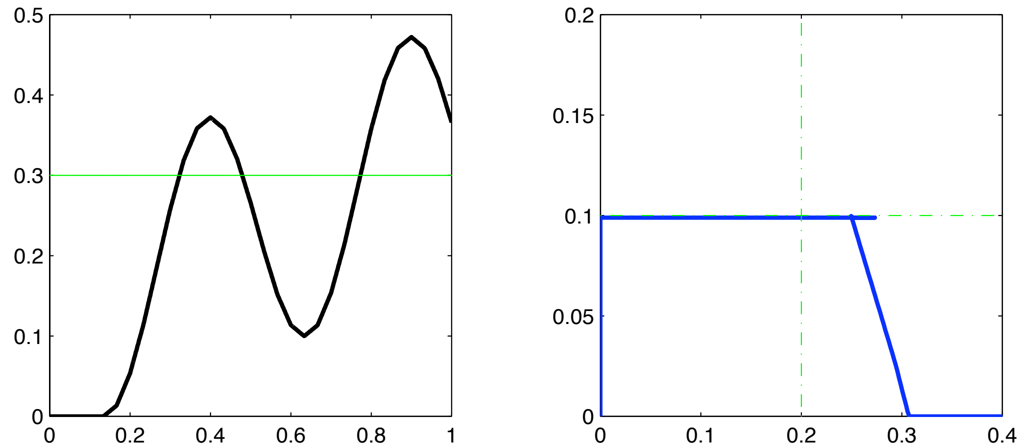

allowing us to conclude that for every $t \leq t^*$, we have that $\bar{L}_t^{(1)}(\mu^*) = 0$. When combined with the fact that $\mu^*(t) = 0$ and, hence, $\bar{L}_t^{(1)}(\mu^*) - \bar{L}_{t^*}^{(1)}(\mu^*) = 0$ for $t > t^*$, this shows that μ^* satisfies condition $(\bar{\mathcal{L}}^*(m) : 1)$. We can write

$$\bar{Q}_t^{(1)}(\mu^*) = \bar{X}_t^{(1)}(\mu^*) = (\bar{Q}_t^P - K_2)^+, \text{ for all } t \leq t^*. \tag{4.4.14}$$

Moreover, for $t \in [0, t^*]$, we have

$$\bar{Q}_t^{(1)}(\mu^*) > K_1 \Leftrightarrow \bar{Q}_t^P > K_1 + K_2. \tag{4.4.15}$$

Focusing on the netput process of the second queue, we obtain

$$\begin{aligned}\bar{X}_t^{(2)}(\mu^*) &= \mathcal{I}_t(\mu^* - \mu^2) \\ &= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P < K_2\}}\, ds \\ &= \int_0^t (\lambda_s - \mu_s^2)\, ds - \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P \geq K_2\}}\, ds.\end{aligned}$$

Using again the integral form (4.4.13) and the definition (4.4.14) along with the the equalities (D.2.4) and (D.2.5) of Appendix D.2, we conclude that

$$\bar{X}_t^{(2)}(\mu^*) = \bar{X}_t^P - (\bar{Q}_t^P - K_2)^+.$$

Next, we evaluate the regulating term for the second queue $\bar{L}^{(2)}(\mu^*)$ on $[0, t^*]$. For every $t \leq t^*$, we have that

$$\begin{aligned}
\bar{L}_t^{(2)}(\mu^*) &= \sup_{s\le t}[-\bar{X}_s^{(2)}(\mu^*)] \\
&= \sup_{s\le t}[-(\bar{X}_s^P - (\bar{Q}_s^P - K_2)^+)] \\
&= \sup_{s\le t}[-\bar{X}_s^P + (\bar{Q}_s^P - K_2)^+] \\
&\le \sup_{s\le t}[-\bar{X}_s^P + \bar{Q}_s^P] \\
&\le \sup_{s\le t}[-\bar{X}_s^P + \bar{X}_s^P + \bar{L}_s^P] \\
&= \bar{L}_t^P.
\end{aligned}$$

We have already proven the opposite inequality in Lemma 4.4.1, so

$$\bar{L}_t^{(2)}(\mu^*) = \bar{L}_t^P, \text{ for all } t \le t^*. \tag{4.4.16}$$

Therefore, for $t \le t^*$, the second queue length equals

$$\bar{Q}_t^{(2)}(\mu^*) = \bar{X}_t^P - (\bar{Q}_t^P - K_2)^+ + \bar{L}_t^P = \bar{Q}_t^P - (\bar{Q}_t^P - K_2)^+ \le K_2. \tag{4.4.17}$$

Our goal of not incurring any penalty in the second queue is, hence, achieved, thus showing that μ^* satisfies condition $(\bar{\mathcal{L}}^*(m) : 3.)$.

Theorem 4.4.7. *Let t^* be defined as in (4.4.10). Then the service discipline μ^* defined by (4.4.9) and (4.4.11) is fluid optimal.*

Proof. Using (4.4.17) and the fact that on $[t^*, T]$ the second queue must be nonincreasing, the performance measure map $\bar{J}$ evaluated at μ^* reduces to

$$\bar{J}(\mu^*) = \text{meas}\,\{s \in [0,T] : \bar{Q}_s^{(1)}(\mu^*) > K_1\}. \tag{4.4.18}$$

Let us consider the following two cases.

Case 1. Let $\bar{Q}_{t^*}^{(1)}(\mu^*) > K_1$. This inequality straightforwardly implies $\mathcal{I}_{t^*}(\lambda) > m + K_1$. Recalling the definition of $\tau : \mathbb{R}_+ \to [0,T]$, given in (3.2.7), we have $t^* \ge \tau(K_1 + m)$. Using the equality (4.4.18), we can write

$$\bar{J}(\mu^*) \le \text{meas}\,\{s \in [0, \tau(K_1+m)] : \bar{Q}_s^{(1)}(\mu^*) > K_1\} + T - \tau(K_1+m).$$

With the fact that $t^* \ge \tau(K_1 + m)$ in mind, we can use the equivalence in (4.4.15) and obtain

$$\begin{aligned}
\bar{J}(\mu^*) &\le \text{meas}\,\{s \in [0, \tau(K_1+m)] : \bar{Q}_s^P > K_1 + K_2\} + T - \tau(K_1+m) \\
&= \bar{J}^P(\tau(K_1+m)) + T - \tau(K_1+m).
\end{aligned}$$

Recalling Corollary 4.4.4, we see that the lower bound stated therein must be attained by policy μ^*.

Case 2. Alternatively, let $\bar{Q}^{(1)}_{t^*}(\mu^*) \le K_1$. Then $\mathcal{I}_{t^*}(\lambda) \le K_1 + m$, and so $t^* \le \tau(K_1 + m)$. For all $t \in (t^*, \tau(K_1 + m))$, by the definition of τ we have $\bar{Q}^{(1)}_t(\mu^*) = \mathcal{I}_t(\lambda) - m \le K_1$. Thus,

$$\begin{aligned}\bar{J}(\mu^*) &\le \text{meas}\,\{s \in [0, \tau(K_1+m)] : \bar{Q}^{(1)}_s(\mu^*) > K_1\} + T - \tau(K_1+m)\\ &= \text{meas}\,\{s \in [0, t^*] : \bar{Q}^{(1)}_s(\mu^*) > K_1\} + T - \tau(K_1+m)\\ &= \bar{J}^P(t^*) + T - \tau(K_1+m),\end{aligned}$$

where the last equality follows from (4.4.15). This, along with Corollary 4.4.4 and the fact that $\bar{J}^P(t^*) \le \bar{J}^P(\tau(K_1+m))$, implies that

$$\bar{J}(\mu^*) = \bar{J}^P(\tau(K_1+m)) + T - \tau(K_1+m).$$

□

4.5 Fluid Limit Analysis in the Finite-Buffer Case.

For any fixed service discipline $\mu \in \mathbb{L}^1_+[0,T]$, again using Theorem 1.2.5, we get

$$\frac{1}{n}X^{(1,n)}(\mu) \to \bar{X}^{(1)}(\mu) = \mathcal{I}(\lambda - \mu),\ a.s.,$$

in the uniform topology. By Corollary 1.2.6, we conclude that also

$$\frac{1}{n}\,Q^{(1,n)}(\mu) \to \bar{Q}^{(1)}(\mu) = \Gamma^{K_1}(\mathcal{I}(\lambda - \mu)),\ a.s., \tag{4.5.1}$$

uniformly, where Γ^{K_1} is the two-sided reflection map on $[0, K_1]$ of Definition 1.2.2. The fluid-limit of the queue length can be rewritten in the standard way as

$$\bar{Q}^{(1)}(\mu) = \bar{X}^{(1)}(\mu) + \bar{L}^{(1)}(\mu) - \bar{U}^{(1)}(\mu), \tag{4.5.2}$$

with $\bar{L}^{(1)}(\mu)$ and $\bar{U}^{(1)}(\mu)$ the regulator maps associated with $\bar{X}^{(1)}(\mu)$ and K_1 as in Definition 1.2.2. As for the second station, using Theorem 1.2.5 and Corollary 1.2.6 again, we have

$$\frac{1}{n}X^{(2,n)}(\mu) \to \bar{X}^{(2)}(\mu) = \mathcal{I}(\mu - \mu^2) - \bar{L}^{(1)}(\mu),\ a.s.,$$

uniformly. Applying once more Corollary 1.2.6, we conclude that

$$\frac{1}{n}\,Q^{(2,n)}(\mu) \to \bar{Q}^{(2)}(\mu) = \mathcal{I}(\mu - \mu^2) - \bar{L}^{(1)}(\mu) + \bar{L}^{(2)}(\mu) - \bar{U}^{(2)}(\mu), \tag{4.5.3}$$

almost surely, in the uniform topology, where $\bar{L}^{(2)}(\mu)$ and $\bar{U}^{(2)}(\mu)$ are the regulator maps associated with $\mathcal{I}(\mu - \mu^2) - \bar{L}^{(1)}(\mu)$ and K_2, as described in Definition 1.2.2.

Having established the limits in (4.5.1) and (4.5.3), we now set up the appropriate fluid-limit analogue of the sequence of performance measures defined in (4.2.3), i.e., we define the mapping $\bar{J}_F : \mathbb{L}^1_+[0,T] \to \mathbb{R}_+$ as

$$\bar{J}_F(\mu) = \bar{U}^{(1)}_T(\mu) + \bar{U}^{(2)}_T(\mu), \text{ for all } \mu \in \mathbb{L}^1_+[0,T].$$

We wish to minimize $\bar{J}_F$ on the set of all admissible service disciplines $\bar{\mathcal{L}}(m)$ and we will do so using a strategy analogous to the one from the infinite buffer case.

Since there are two regulator processes involved in both stations, the notation is rather cumbersome in this case. Also certain subclasses of admissible controls appear to be clearly suboptimal. In particular, it seems reasonable for the controller to use only the service disciplines μ conforming to the following properties.

- $\bar{L}^{(1)}_T(\mu) = 0$.
  In both the infinite-buffer and the finite-buffer problems it is in the controller's best interest not to serve "in vain" the first station, as he/she might run out of fuel before time T and then possibly incur unnecessary cost in the first station. Hence, the condition indicated above is indeed sensible.

- $\bar{U}^{(2)}_T(\mu) = 0$.
  In the infinite-buffer case, the penalties in both queues were accumulated independently when the queues transgressed over their thresholds. It was not important which queue was the one causing an increase in the penalty as the costs were equal for both queues and all the jobs remained in the queues until they were served or time T was reached. To the contrary, in the finite-buffer case it is important that the first station precedes the second station. Specifically, it is possible for the following scenario to take place:

  1. A job arrives in the first station, while the queue is at its full capacity.
  2. The controller decides to speed up the service in order not to lose the job that arrived.
  3. The job gets served in the first station and moves on to the second one.
  4. The second station is at its full capacity and that same job does not get completed, after all.
  5. The epilogue is that the controller still got penalized for losing the job from the second station, while he "wasted fuel" on that same job in the first station.

  We need to remember that the controller is only penalized for the jobs lost - he/she does not get rewarded based on how many jobs leave the system in the allotted time. Hence, it is not necessarily in the controller's best interest to be serving in the first station relying on the second station to actually complete the jobs that get sent to it. This explains why it is preferable to lose jobs in the first station as opposed to losing jobs in the second station, which is exactly enforced in the announced condition.

Let us formalize the rationale behind the above conditions by proving it suffices to optimize the performance measure $\bar{J}_F$ over the admissible service disciplines satisfying them.

Lemma 4.5.1. *Let μ be an admissible service discipline. Then there exists an admissible discipline μ'' which satisfies the following three conditions:*

(i) $\bar{L}_T^{(1)}(\mu'') = 0$.

(ii) $\bar{U}_T^{(2)}(\mu'') = 0$.

(iii) $\bar{J}_F(\mu) \geq \bar{J}_F(\mu'')$.

Proof. Let μ be an admissible service discipline. The strategy of the proof consists of constructing from μ the new service discipline μ'' in two steps. First, we define the auxiliary function μ' satisfying conditions (i) and (iii) in the lemma. This will simply ease the exposition of the next step. Then, we use μ' to find a μ'' satisfying all three conditions posited in the lemma.

We define the subset A of the interval $[0,T]$ as $A = \{t \in [0,T] : \bar{Q}_t^{(1)}(\mu) = 0\}$ and the new service discipline μ' as

$$\mu' = \mu \mathbf{1}_{A^c} + (\lambda \wedge \mu)\mathbf{1}_A.$$

Alternatively, we can describe μ' through the value of its integral. For every t, we have

$$\begin{aligned}
\mathcal{I}_t(\mu') &= \int_0^t \mathbf{1}_A(s)(\mu_s \wedge \lambda_s)\,ds + \int_0^t \mathbf{1}_{A^c}(s)\mu_s\,ds \\
&= \int_0^t \mathbf{1}_A(s)\mu_s \mathbf{1}_{\{\mu_s \leq \lambda_s\}}\,ds + \int_0^t \mathbf{1}_A(s)\lambda_s \mathbf{1}_{\{\mu_s > \lambda_s\}}\,ds + \int_0^t \mathbf{1}_{A^c}(s)\mu_s\,ds \\
&= \int_0^t \mu_s\,ds - \int_0^t \mathbf{1}_A(s)(-\lambda_s + \mu_s)^+\,ds \\
&= \mathcal{I}_t(\mu) - \bar{L}_t^{(1)}(\mu).
\end{aligned} \tag{4.5.4}$$

We claim (and show below) that μ' is itself an admissible service discipline which does not require lower regulation in the first station and does not perform worse that μ.

Admissibility. All service disciplines in the fluid-limit context are assumed to be deterministic. So it suffices to verify that μ' satisfies the constraint m on the cumulative service available.

$$\mathcal{I}_T(\mu') = \int_{A^c} \mu_s\,ds + \int_A \lambda_s \wedge \mu_s\,ds \leq \int_0^T \mu_s\,ds \leq m.$$

Absence of lower regulation. The netput process in the first station generated by the service discipline μ' is

$$\bar{X}^{(1)}(\mu') = \mathcal{I}(\lambda) - \mathcal{I}(\mu').$$

Using (4.5.4), we see that the last expression equals

$$\bar{X}^{(1)}(\mu') = \mathcal{I}(\lambda) - \mathcal{I}(\mu) + \bar{L}^{(1)}(\mu). \tag{4.5.5}$$

On the other hand, the queue length associated with the service discipline μ is

$$\bar{Q}^{(1)}(\mu) = \mathcal{I}(\lambda) - \mathcal{I}(\mu) + \bar{L}^{(1)}(\mu) - \bar{U}^{(1)}(\mu)$$

which, thanks to (4.5.5), can be rewritten as

$$\bar{Q}^{(1)}(\mu) = \bar{X}^{(1)}(\mu') - \bar{U}^{(1)}(\mu). \tag{4.5.6}$$

The processes $\bar{L}^{(1)}(\mu')$ and $\bar{U}^{(1)}(\mu')$ are, by Proposition 1.2.3, the minimal nondecreasing processes rendering the queue length to be between 0 and K_1, so in view of (4.5.6) we conclude that

$$\begin{aligned} \bar{L}^{(1)}(\mu') &\equiv 0, \\ \bar{U}^{(1)}(\mu') &= \bar{U}^{(1)}(\mu). \end{aligned} \tag{4.5.7}$$

Performance. The netput process for the second queue generated by the service discipline μ' is

$$\bar{X}^{(2)}(\mu') = \mathcal{I}(\mu') - \mathcal{I}(\mu^2). \tag{4.5.8}$$

Plugging (4.5.4) into the last expression, we obtain

$$\bar{X}^{(2)}(\mu') = \mathcal{I}(\mu) - \bar{L}^{(1)}(\mu) - \mathcal{I}(\mu^2). \tag{4.5.9}$$

We see that this is the same process as the netput process in the second station associated with μ. Hence, all the derived processes (queue length, regulator processes, etc.) must be the same as well. In particular,

$$\bar{U}^{(2)}_T(\mu') = \bar{U}^{(2)}_T(\mu).$$

Gathering together the second line of (4.5.7) and the last display, we obtain that $\bar{J}_F(\mu') = \bar{J}_F(\mu)$.

Having established the desired properties of μ', we move on to the construction of μ''. We define $B = \{t \in [0,T] : \bar{Q}^{(2)}_t(\mu') = K_2\}$ and the integrable function

$$\mu'' = \mu' \mathbf{1}_{B^c} + (\mu' \wedge \mu^2)\mathbf{1}_B.$$

A useful representation of the integral of μ'' is the following:

$$\begin{aligned} \mathcal{I}_t(\mu'') &= \int_0^t \mu'_s \mathbf{1}_{B^c}(s)\, ds + \int_0^t (\mu'_s \wedge \mu^2_s)\mathbf{1}_B(s)\, ds \\ &= \int_0^t \mu'_s\, ds + \int_0^t (\mu'_s \wedge \mu^2_s - \mu'_s)\mathbf{1}_B(s)\, ds \\ &= \mathcal{I}_t(\mu') - \bar{U}^{(2)}_t(\mu'), \text{ for every } t. \end{aligned} \tag{4.5.10}$$

Now we prove that μ'' indeed satisfies the three conditions declared in the lemma.

Admissibility. The appropriate bound on the integral of the function μ'' over the interval $[0,T]$ is established as follows.

$$\mathcal{I}_T(\mu'') = \int_{B^c} \mu'_s\, ds + \int_B \mu'_s \wedge \mu^2_s\, ds \leq \int_0^T \mu'_s\, ds \leq m.$$

Absence of upper regulation in the second station. The netput process for the second station for the service discipline μ'' is simply

$$\bar{X}^{(2)}(\mu'') = \mathcal{I}(\mu'') - \bar{L}^{(1)}(\mu'') - \mathcal{I}(\mu^2).$$

Inserting the result of (4.5.10) into the expression in the last display, we get

$$\bar{X}^{(2)}(\mu'') = \mathcal{I}(\mu') - \bar{U}^{(2)}(\mu') - \bar{L}^{(1)}(\mu'') - \mathcal{I}(\mu^2). \tag{4.5.11}$$

Applying the two-sided regulator map to the netput process $\bar{X}^{(2)}(\mu'')$ and expanding the result utilizing the minimal lower and upper regulators in the sense of Definition 1.2.2, we obtain

$$\bar{Q}^{(2)}(\mu'') = \bar{X}^{(2)}(\mu'') + \bar{L}^{(2)}(\mu'') - \bar{U}^{(2)}(\mu''). \tag{4.5.12}$$

Recalling the netput process $\bar{X}^{(2)}(\mu')$ in (4.5.8), we see that the result of the regulation in the second queue when μ' is used can be expressed as

$$\begin{aligned}\bar{Q}^{(2)}(\mu') &= \mathcal{I}(\mu') - \mathcal{I}(\mu^2) + \bar{L}^{(2)}(\mu') - \bar{U}^{(2)}(\mu') \\ &= \mathcal{I}(\mu') - \bar{U}^{(2)}(\mu') - \bar{L}^{(1)}(\mu'') - \mathcal{I}(\mu^2) + \bar{L}^{(1)}(\mu'') + \bar{L}^{(2)}(\mu').\end{aligned}$$

According to (4.5.11) we can rewrite the the last result as

$$\bar{Q}^{(2)}(\mu') = \bar{X}^{(2)}(\mu'') + \bar{L}^{(1)}(\mu'') + \bar{L}^{(2)}(\mu'). \tag{4.5.13}$$

By definition, the quantity in the last display is always nonnegative and smaller than or equal to K_2. Comparing (4.5.12) to (4.5.13) and with Proposition 1.2.3 in mind, we can conclude that $\bar{U}^{(2)}(\mu'') \leq 0$ and hence

$$\bar{U}^{(2)}(\mu'') \equiv 0. \tag{4.5.14}$$

Absence of lower regulation in the first station. The netput process for the first queue when μ'' is used equals

$$\bar{X}^{(1)}(\mu'') = \mathcal{I}(\lambda) - \mathcal{I}(\mu'') = \mathcal{I}(\lambda) - \mathcal{I}(\mu') + \bar{U}^{(2)}(\mu').$$

Therefore, the queue length obtained when the two-sided regulator map is applied to $\bar{X}^{(1)}(\mu'')$ is given by

$$\bar{Q}^{(1)}(\mu'') = \mathcal{I}(\lambda) - \mathcal{I}(\mu') + \bar{U}^{(2)}(\mu') + \bar{L}^{(1)}(\mu'') - \bar{U}^{(1)}(\mu''),$$

where $\bar{L}^{(1)}(\mu'')$ and $\bar{U}^{(1)}(\mu'')$ are the regulator maps associated with $\bar{X}^{(1)}(\mu'')$ and K_1 in the sense of Definition 1.2.2. Simultaneously, the queue length in the first station produced by the service discipline μ' is equal to

$$\begin{aligned}\bar{Q}^{(1)}(\mu') &= \mathcal{I}(\lambda) - \mathcal{I}(\mu') + \bar{L}^{(1)}(\mu') - \bar{U}^{(1)}(\mu') \\ &= \mathcal{I}(\lambda) - \mathcal{I}(\mu') + \bar{U}^{(2)}(\mu') + \bar{L}^{(1)}(\mu') - \bar{U}^{(1)}(\mu') - \bar{U}^{(2)}(\mu').\end{aligned}$$

The last display, along with Proposition 1.2.3, implies that

$$\begin{aligned}\bar{L}^{(1)}(\mu'') &\leq \bar{L}^{(1)}(\mu') \\ \bar{U}^{(1)}(\mu'') &\leq \bar{U}^{(1)}(\mu') + \bar{U}^{(2)}(\mu').\end{aligned} \tag{4.5.15}$$

Using (4.5.7) and the first line above, we conclude that $\bar{L}^{(1)}(\mu'') \equiv 0$.

Performance. Combining (4.5.14) with the second line of (4.5.15), we get that μ'' does indeed outperform μ' and, hence, μ. This is the last claim that was to be proven. □

From now on, we shall refer to the subclass of the space of all admissible disciplines $\bar{\mathcal{L}}(m)$ satisfying conditions (i) and (ii) of Lemma 4.5.1 as $\bar{\mathcal{L}}''(m)$.

4.5.1 The Pooled Queue

Next, we consider the finite-buffer counterpart of the pooled queue introduced in (4.4.2). The idea of "pooling" the system by ignoring the middle station is the same, except that in the present case the pooled queue has a finite buffer of size $K = K_1 + K_2$. The queue length in this system is then

$$\bar{Q}^{(P)} = \Gamma^K(\mathcal{I}(\lambda - \mu^2)).$$

As usual, it is convenient to rewrite this process as

$$\bar{Q}^{(P)} = \mathcal{I}(\lambda - \mu^2) + \bar{L}^P - \bar{U}^P,$$

with $\bar{L}^P$ and $\bar{U}^P$ being the regulators associated with $\mathcal{I}(\lambda - \mu^2)$ and K, as in Definition 1.2.2.

As we have witnessed in Subsection 4.4.2, the pooled queue is a useful instrument for the analysis of the corresponding tandem system only until the time when either the service discipline has reached the imposed constraint or, even if the total amount of service used did not reach the constraint m yet, there is not enough time remaining to ever bring the first queue under its threshold. We again observe the entire evolution of the performance of the pooled queue in comparison with the performances of the tandem system with varying controls.

Based on the findings of Lemma 4.5.1, it suffices to consider the performance of the service disciplines in $\bar{\mathcal{L}}''(m)$ in order to establish a lower bound on the performance of all admissible service disciplines.

Lemma 4.5.2. *For every $\mu \in \bar{\mathcal{L}}''(m)$ the following inequalities hold true:*

(i) $\bar{U}^P \leq \bar{U}^{(1)}(\mu)$;

(ii) $\bar{L}^P \leq \bar{L}^{(2)}(\mu)$.

Proof. We start by defining two sequences of deterministic times in $[0,T]$ as follows:

$$\begin{aligned} t_1 &= \inf\{t > 0 : \bar{Q}^P_t = K\} \wedge T, \\ s_i &= \inf\{t > t_i : \bar{Q}^P_t = 0\} \wedge T, \text{ for } i \geq 1, \\ t_i &= \inf\{t > s_{i-1} : \bar{Q}^P_t = K\} \wedge T, \text{ for } i > 1. \end{aligned} \tag{4.5.16}$$

Let us fix an arbitrary admissible $\mu \in \mathcal{L}''(m)$. Then we have that

$$\mathcal{I}(\lambda) \geq \mathcal{I}(\mu) + \bar{U}^{(1)}(\mu) \geq \mathcal{I}(\mu), \tag{4.5.17}$$

where the first inequality uses the facts that $\bar{L}^{(1)}(\mu) \geq 0$ and $\bar{Q}^{(1)}(\mu) \geq 0$, while the second inequality hold trivially because $\bar{U}^{(1)}(\mu) \geq 0$.

We will prove both claims in the lemma simultaneously, using the principle of mathematical induction. On the segment $[0, t_1]$, the length of the pooled queue equals

$$\bar{Q}^P = \bar{X}^P + \bar{L}^P.$$

Its lower regulator $\bar{L}^P$ can be written, for all $t \in [0, t_1]$, as

$$\bar{L}^P_t = \sup_{s \leq t}[-\mathcal{I}(\lambda - \mu^2)].$$

Using the last equality and the inequality (4.5.17), we get

$$\bar{L}^P_t \leq \sup_{s \leq t}[-\mathcal{I}(\mu - \mu^2)] = \bar{L}^{(2)}_t(\mu),$$

where the latter equality holds since $\mu \in \bar{\mathcal{L}}''(m)$ implies that $\bar{U}^{(2)}(\mu) = 0$. On the other hand, for every $t \leq t_1$, we have that $\bar{U}^P_t = 0$. By the nonnegativity of the regulator maps, the first announced inequality holds as well for all $t \in [0, t_1]$.

Next, we focus on the segment $[t_1, s_1]$. By definition, there is no need for lower regulation of the pooled queue in this region. So, we have that for every $t \in [t_1, s_1]$,

$$\bar{L}^P_t = \bar{L}^P_{t_1} \tag{4.5.18}$$

and, therefore, $\bar{Q}^P_t = \bar{X}^P_t + \bar{L}^P_{t_1} - \bar{U}^P_t$.

Since the regulator maps are by definition nondecreasing, equality (4.5.18) and the just proven validity of the posited inequalities on the segment $[0, t_1]$ imply that

$$\bar{L}^P_t \leq \bar{L}^{(2)}_{t_1}(\mu) \leq \bar{L}^{(2)}_t(\mu), \text{ for every } t \in [t_1, s_1]. \tag{4.5.19}$$

As for the inequality involving the upper regulators, we have that

$$\bar{U}^P_t = \sup_{s \leq t}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu^2) + \bar{L}^P_s - K]^+, \text{ for every } t \in [t_1, s_1].$$

Since the process $\bar{U}^P$ can only increase when $\bar{Q}^P = K$, we can rewrite the last equality using (4.5.18) and (4.5.19) as

$$\begin{aligned}
\bar{U}^P_t &= \sup_{t_1 \leq s \leq t}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu^2) + \bar{L}^P_{t_1} - K]^+ \\
&\leq \sup_{t_1 \leq s \leq t}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu) - K_1]^+ + \sup_{t_1 \leq s \leq t}[\mathcal{I}_s(\mu) - \mathcal{I}_s(\mu^2) + \bar{L}^P_{t_1} - K_2]^+ \\
&\leq \sup_{s \leq t}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu) - K_1]^+ + \sup_{t_1 \leq s \leq t}[\bar{Q}^{(2)}_s(\mu) - K_2]^+ \\
&= \sup_{s \leq t}[\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu) - K_1]^+ \\
&= \bar{U}^{(1)}_t(\mu), \text{ for every } t \in [t_1, s_1].
\end{aligned}$$

Let us assume that the posited inequalities hold true on the entire region $[0, s_{i-1}]$, for some $i \geq 2$. We will next prove that those claims necessarily carry over from the stated inductive hypothesis to the segment $[s_{i-1}, s_i]$. For all $t \in [s_{i-1}, t_i]$ the pooled queue is strictly below the level K. So

$$\bar{U}^P_t = \bar{U}^P_{s_{i-1}}, \text{ for every } t \in [s_{i-1}, t_i]. \tag{4.5.20}$$

Hence, $\bar{Q}^P_t = \bar{X}^P_t + \bar{L}^P_t - \bar{U}^P_{s_{i-1}}$, and

$$\bar{L}^P_t = \bar{L}^P_{s_{i-1}} \vee \sup_{s_{i-1} \leq s \leq t_i}[-\mathcal{I}_s(\lambda) + \mathcal{I}_s(\mu^2) + \bar{U}^P_{s_{i-1}}]^+. \tag{4.5.21}$$

Using the inductive hypothesis, we have that

$$\bar{L}^P_{s_{i-1}} \leq \bar{L}^{(2)}_{s_{i-1}}(\mu). \tag{4.5.22}$$

On the other hand, using (4.5.17), we obtain

$$\sup_{s_{i-1} \leq s \leq t_i}[-\mathcal{I}_s(\lambda) + \mathcal{I}_s(\mu^2) + \bar{U}^P_{s_{i-1}}]^+ = \sup_{s_{i-1} \leq s \leq t_i}[-\mathcal{I}_s(\mu) - \bar{U}^{(1)}_s(\mu) + \mathcal{I}_s(\mu^2) + \bar{U}^P_{s_{i-1}}]^+.$$

Thanks to the fact that regulator maps are nondecreasing and to the inductive hypothesis, the quantity in the above expression is exceeded by

$$\sup_{s_{i-1}\le s\le t_i}[-\mathcal{I}_s(\mu)+\mathcal{I}_s(\mu^2)]^+.$$

Combining the last upper bound with the inequality (4.5.22) and the expansion in (4.5.21), we obtain the desired inequality.

The inequality involving the upper regulators is again a simple consequence of the monotonicity of $\bar{U}^{(1)}$ and (4.5.20).

Finally, we tackle the segment $[t_i, s_i]$. Here, there is no possibility of upward pushing in the pooled queue, so $\bar{L}^P_t = \bar{L}^P_{t_i}$, for every $t \in [t_i, s_i]$. We can immediately conclude from the monotonicity of the lower regulator in the second queue that the second proposed inequality holds true in this region. For every $t \in [t_i, s_i]$, the amount of downward pushing in the pooled queue during the segment $[t_i, t]$ is given by

$$\begin{aligned}\sup_{t_i\le s\le t}&[\mathcal{I}_s(\lambda)-\mathcal{I}_s(\mu^2)+\bar{L}^P_{t_i}-K]^+\\&\le \sup_{t_i\le s\le t}[\mathcal{I}_s(\lambda)-\mathcal{I}_s(\mu)-K_1]^+ + \sup_{t_i\le s\le t}[\mathcal{I}_s(\mu)-\mathcal{I}_s(\mu^2)+\bar{L}^P_{t_i}-K_2]^+.\end{aligned}\tag{4.5.23}$$

Thanks to the inductive hypothesis, the second term on the right-hand side of the above equation is bounded from above by

$$\sup_{t_i\le s\le t}[\mathcal{I}_s(\mu)-\mathcal{I}_s(\mu^2)+\bar{L}^{(2)}_s(\mu)-K_2]^+ = \sup_{t_i\le s\le t}[\bar{Q}^{(2)}_s(\mu)-K_2]^+ = 0.$$

We can write

$$\bar{U}^P_t = \bar{U}^P_{t_i} \vee \sup_{t_i\le s\le t}[\mathcal{I}_s(\lambda)-\mathcal{I}_s(\mu^2)+\bar{L}^P_{t_i}-K]^+.$$

According to (4.5.23), the validity of the first proposed claim on the segment $[s_{i-1}, t_i]$, and the last equality, we have

$$\bar{U}^P_t \le \bar{U}^{(1)}_{t_i}(\mu) \vee \sup_{t_i\le s\le t}[\mathcal{I}_s(\lambda)-\mathcal{I}_s(\mu)-K_1]^+ = \bar{U}^{(1)}_t(\mu), \text{ for every } t\in[t_i,s_i].$$

□

We proceed with the introduction of a control-dependent time instance $\tau^* : \mathbb{L}^1_+[0,T] \to [0,T]$ given by

$$\tau^*(\mu) = \inf\{t\in[0,T] : \mathcal{I}_t(\mu) = m\} \wedge T, \text{ for all } \mu \in \mathbb{L}^1_+[0,T].$$

In words, the function τ^* returns for any service discipline the first time its cumulative service reaches the level m, i.e., the time at which the service constraint we imposed on admissible disciplines is reached and there can no longer be any service rendered in the first station.

We claim that the service discipline $\mu^{*,F}$ we define next is fluid-optimal in the context of the performance measure $\bar{J}_F$. First, let the auxiliary function $\mu^{aux,F} \in \mathbb{L}^1_+[0,T]$ be given for every t as

$$\mu_t^{aux,F} = \begin{cases} \lambda_t & \text{if } \bar{Q}_t^P < K_2 \\ \mu_t^2 & \text{if } \bar{Q}_t^P \geq K_2. \end{cases}$$

Then, the instant at which the service discipline $\mu^{aux,F}$ reaches the constraint imposed on the total amount of service available is $\tau^*(\mu^{aux,F})$. The admissible service discipline which behaves like $\mu^{aux,F}$ for as long as there is service available is, hence,

$$\mu^{*,F} = \mu^{aux,F}\mathbf{1}_{[0,\tau^*(\mu^{aux,F})]}. \tag{4.5.24}$$

Clearly, we have that $\tau^*(\mu^{aux,F}) = \tau^*(\mu^{*,F})$. For typographical reasons, we denote this instant by $T^* = \tau^*(\mu^{*,F})$.

By construction, $\mu^{*,F}$ is admissible. We claim that it, moreover, satisfies equalities (i) and (ii) of Lemma 4.5.1.

Lemma 4.5.3. *The service discipline $\mu^{*,F}$ belongs to the class $\bar{\mathcal{L}}''(m)$. Additionally, we have that*

$$\bar{U}_t^{(1)}(\mu^{*,F}) = \bar{U}_t^P, \text{ for every } t \leq T^*. \tag{4.5.25}$$

Proof. We first justify the validity of the two equalities stated in Lemma 4.5.1. In order to do this, it suffices to verify that the equalities hold true at time T^*. From that point on there can no longer be any service in the first station which prevents the increase of both the lower regulator for the first queue and the upper regulator for the second queue.

Let us begin with the fact that the lower regulator in the first queue vanishes for $\mu^{*,F}$. For every $t \leq T^*$, we have that

$$\begin{aligned} \bar{X}_t^{(1)}(\mu^{*,F}) &= \int_0^t (\lambda_s - \mu_s^{*,F})\,ds \\ &= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P \geq K_2\}}\,ds \\ &= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{K > \bar{Q}_s^P \geq K_2\}}\,ds + \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P = K\}}\,ds \\ &= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{K > \bar{Q}_s^P \geq K_2\}}\,ds + \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P = K\}}\,ds \\ &\quad - \int_0^t (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P = K\}}\,ds + \int_0^t (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P = K\}}\,ds. \end{aligned} \tag{4.5.26}$$

Using the results (D.2.3), (D.2.4) and (D.2.5) of Appendix D.2 with

$$f_s := (\lambda_s - \mu_s^2)\mathbf{1}_{\{0<\bar{Q}_s^P<K\}} + (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P=0\}} - (\lambda_s - \mu_s^2)^-\mathbf{1}_{\{\bar{Q}_s^P=K\}}, \text{ for all } s \in [0,T] \tag{4.5.27}$$

and the integral form for the length of the pooled queue

$$\bar{Q}_t^{(P)} = \int_0^t f_s\, ds, \tag{4.5.28}$$

we transform the final sum of integrals in (4.5.26) to obtain

$$\bar{X}_t^{(1)}(\mu^{*,F}) = (\bar{Q}_t^P - K_2)^+ + \bar{U}_t^P. \tag{4.5.29}$$

An immediate consequence of the above calculation is that

$$\bar{X}_t^{(1)}(\mu^{*,F}) - \bar{U}_t^P = (\bar{Q}_t^P - K_2)^+, \text{ for every } t \leq T^*. \tag{4.5.30}$$

The right-hand side is bounded between 0 and K_1. Therefore, by Proposition 1.2.3, we have that

$$\bar{L}^{(1)}(\mu^{*,F}) \equiv 0 \tag{4.5.31}$$

and

$$\bar{U}_t^{(1)}(\mu^{*,F}) \leq \bar{U}_t^P, \text{ for every } t \leq T^*. \tag{4.5.32}$$

The steps in the proof for absence of upper regulation in the second queue are similar. Based on the fact that there is no lower regulation in the first station, the netput process in the second station equals for every $t \leq T^*$

$$\begin{aligned}
\bar{X}_t^{(2)}(\mu^{*,F}) &= \int_0^t (\mu_s^{*,F} - \mu_s^2)\, ds \\
&= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P<K_2\}}\, ds \\
&= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{0<\bar{Q}_s^P<K_2\}}\, ds + \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds \\
&= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{0<\bar{Q}_s^P<K_2\}}\, ds + \int_0^t (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds \\
&\quad - \int_0^t (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds + \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds \\
&= \int_0^t (\lambda_s - \mu_s^2)\mathbf{1}_{\{0<\bar{Q}_s^P<K_2\}}\, ds + \int_0^t (\lambda_s - \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds \\
&\quad - \int_0^t (-\lambda_s + \mu_s^2)^+\mathbf{1}_{\{\bar{Q}_s^P=0\}}\, ds.
\end{aligned} \tag{4.5.33}$$

Simplifying the sum of the first two integrals in the result above using (D.2.3), (D.2.4) and (D.2.5) in Appendix D.2 and the integral form of the length of the pooled queue in (4.5.28), we obtain

$$\bar{X}_t^{(2)}(\mu^{*,F}) = \bar{Q}_t^P \wedge K_2 - \bar{L}_t^P, \text{ for every } t \leq T^*. \tag{4.5.34}$$

Somewhat rearranging the terms in the above expression, we get

$$\bar{X}_t^{(2)}(\mu^{*,F}) + \bar{L}_t^P = \bar{Q}_t^P \wedge K_2, \text{ for every } t \leq T^*.$$

The right-hand side above is nonnegative and restricted to be below or at the level K_2. By Proposition 1.2.3, we conclude that

$$\bar{U}^{(2)}(\mu^{*,F}) \equiv 0. \tag{4.5.35}$$

Therefore, $\mu^{*,F} \in \mathcal{L}''(m)$. By the bounds established in Lemma 4.5.2, we conclude that $\bar{U}_t^{(1)}(\mu^{*,F}) = \bar{U}_t^P$, for every $t \leq T^*$. □

Note that Lemma 4.5.3 shows that the lower bound on the performance measure established in Lemma 4.5.2 is indeed tight, as it is attained by the service discipline $\mu^{*,F}$.

Proposition 4.5.4. *The service discipline $\mu^{*,F}$ is fluid-optimal for $\bar{J}_F$. In other words, for any other admissible service discipline μ, we have that $\bar{J}_F(\mu) \geq \bar{J}_F(\mu^{*,F})$.*

Proof. Clearly, it suffices to compare the performance of $\mu^{*,F}$ to the performance of other elements of the class $\bar{\mathcal{L}}''(m)$. So, let us temporarily fix some $\mu \in \bar{\mathcal{L}}''(m)$. By definition, we have that $\bar{U}_T^{(2)}(\mu) = 0$, and, hence, that $\bar{J}_F(\mu) = \bar{U}_T^{(1)}(\mu)$.

The total amount of downward pushing in the first queue is evaluated as

$$\bar{U}_T^{(1)}(\mu) = \bar{U}_{T^*}^{(1)}(\mu) \vee \sup_{s \in (T^*,T]} [\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu) + \bar{L}_s^{(1)}(\mu) - K_1]^+.$$

Since $\mu \in \bar{\mathcal{L}}''(m)$, this can be rewritten as

$$\bar{U}_T^{(1)}(\mu) = \bar{U}_{T^*}^{(1)}(\mu) \vee \sup_{s \in (T^*,T]} [\mathcal{I}_s(\lambda) - \mathcal{I}_s(\mu) - K_1]^+.$$

Using Lemma 4.5.2 and the constraint on $\mathcal{I}_T(\mu)$, the last expression gets transformed into

$$\bar{U}_T^{(1)}(\mu) \geq \bar{U}_{T^*}^P \vee \sup_{s \in (T^*,T]} [\mathcal{I}_s(\lambda) - m - K_1]^+ = \bar{U}_{T^*}^P \vee [\mathcal{I}_T(\lambda) - m - K_1]^+. \tag{4.5.36}$$

On the other hand, at $\mu^{*,F}$ the total amount of upper regulation in the first queue reads

$$\bar{U}_T^{(1)}(\mu^{*,F}) = \bar{U}_{T^*}^{(1)}(\mu^{*,F}) \vee \sup_{s \in (T^*,T]} [\mathcal{I}_s(\lambda) - m - K_1]^+.$$

Using Lemma 4.5.3 and the definition of T^*, we rewrite the last equality as

$$\bar{U}_T^{(1)}(\mu^{*,F}) = \bar{U}_{T^*}^P \vee [\mathcal{I}_T(\lambda) - m - K_1]^+. \tag{4.5.37}$$

Comparing expressions (4.5.36) and (4.5.37), we obtain the desired inequality. □

Remark 4.5.1. This is a good time to discuss the rationale behind the introduction of the auxiliary infinite-buffer problem. The problem featuring only one-sided regulator maps is evidently more tractable. It also provided us with a class of fluid-optimal class that can be easily described. Using this class as a starting point allows us to analyze the performance of the finite-buffer system on this class, which, in turn, helps us to infer the optimality conditions in the more complicated finite-buffer case.

4.6 Asymptotic Optimality - Infinite Buffers

Definition 4.6.1. A sequence of service processes $\{\mu_n\}$ is called *admissible* if $\mu_n \in \mathcal{L}^{(n)}(m)$ for all n. Furthermore, an admissible sequence $\{\mu_n^*\}$ is called *asymptotically optimal* for the sequence of performance measures $\{J^{(n)}\}$, if

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}(\mu_n^*)] \geq 0,$$

for any other admissible sequence $\{\mu_n\}$.

The following is a straightforward consequence of the above definition.

Lemma 4.6.2. *If $\{\mu_n\}$ and $\{\mu_n'\}$ are two asymptotically optimal sequences, then*

$$\lim_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}(\mu_n')] = 0.$$

It is our goal to find an asymptotically optimal sequence. The definition of asymptotic optimality does not allow for an easy algorithm for either verifying if a given sequence is asymptotically optimal or constructing an asymptotically optimal sequence. We propose the following, more operational, criterion for asymptotic optimality.

Proposition 4.6.3. *Let $\{\tilde{J}_{LB}^{(n)}\}$ be a sequence of random variables such that, for any admissible sequence $\{\mu_n\}$*

$$J^{(n)}(\mu_n) \geq \tilde{J}_{LB}^{(n)}, \text{a.s., for every } n. \tag{4.6.1}$$

Furthermore, let the admissible sequence $\{\tilde{\mu}_n\}$ satisfy

$$\lim_{n\to\infty} \mathbb{E}[J^{(n)}(\tilde{\mu}_n) - \tilde{J}_{LB}^{(n)}] = 0. \tag{4.6.2}$$

Then the sequence $\{\tilde{\mu}_n\}$ is asymptotically optimal in the sense of Definition 4.6.1.

Proof. Let $\{\tilde{\mu}_n\}$ and $\{\tilde{J}^{(n)}_{LB}\}$ satisfy the conditions of the proposition, and let $\{\mu_n\}$ be an arbitrary sequence of admissible service disciplines. Using assumption (4.6.1), we get

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - \tilde{J}^{(n)}_{LB}] \geq 0,$$

which combined with (4.6.2) implies

$$\begin{aligned}&\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}(\tilde{\mu}_n)] \\ &\qquad \geq \liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - \tilde{J}^{(n)}_{LB}] + \lim_{n\to\infty} \mathbb{E}[\tilde{J}^{(n)}_{LB} - J^{(n)}(\tilde{\mu}_n)] \geq 0.\end{aligned}$$

□

The following is a trivial consequence of Lemma 4.6.2 and Proposition 4.6.3.

Corollary 4.6.4. *Suppose that a sequence of random variables $\{\tilde{J}^{(n)}_{LB}\}$ and an asymptotically optimal sequence $\{\tilde{\mu}_n\}$ are given as in Proposition 4.6.3. Then for all asymptotically optimal sequences $\{\mu_n\}$ we have that*

$$\lim_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - \tilde{J}^{(n)}_{LB}] = 0.$$

In the next subsection, we will analyze the pooled queue in order to identify an appropriate sequence $\{\tilde{J}^{(n)}_{LB}\}$ of random variables of Corollary 4.6.4.

4.6.1 Analysis of the Pooled Queues

Let us define the pooled queues - queues whose arrival rates correspond to the arrival rates in the tandem and whose service rates correspond to those in the second station in the tandem. Formally, for all n we introduce the queue-length process

$$Q^{(P,n)} = N_1^+(n\mathcal{I}(\lambda)) - N_2^-(n\mathcal{I}(\mu^2)) + L^{(P,n)}, \tag{4.6.3}$$

with $L^{(P,n)}$ being the correction term arising from the one-sided regulator of Definition 1.2.1 applied to the process $N_1^+(n\mathcal{I}(\lambda)) - N_2^-(n\mathcal{I}(\mu^2))$. The following is a trivial consequence of Corollary 1.2.6.

Corollary 4.6.5. *As $n \to \infty$, $\frac{1}{n}Q^{(P,n)} \to \bar{Q}^P$, a.s., uniformly, with $Q^{(P,n)}$ from (4.6.3) and $\bar{Q}^P$ from (4.4.2).*

The behavior of $Q^{(P,n)}$ in the vicinity of $K = K_1 + K_2$ will be of particular interest to us. The first lemma allows us to restrict our attention to the limiting behavior of the processes $Q^{(P,n)}$ on a particular "grid". Note that its claim trivially holds true for integer-valued K_1 and K_2.

Lemma 4.6.6.

$$\lim_{n\to\infty} \mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(\lfloor nK_1\rfloor + \lfloor nK_2\rfloor) < \frac{1}{n}Q_t^{(P,n)} \leq K_1 + K_2\right\}\right] = 0.$$

Proof. In order to ease the exposition, let us introduce a sequence of constants $\{k(n)\}$ given by $k(n) := \frac{1}{n}(\lfloor nK_1\rfloor + \lfloor nK_2\rfloor)$. Also recall that $K = K_1 + K_2$.

For any fixed n, we can partition the set in the expression above and transform the expectation into a sum of expectations in the following way:

$$\begin{aligned}
&\mathbb{E}\left[meas\left\{t \in [0,T] : k(n) < \frac{1}{n}Q_t^{(P,n)} \leq K\right\}\right] \\
&= \mathbb{E}\left[meas\left\{t \in [0,T] : k(n) < \frac{1}{n}Q_t^{(P,n)} \leq K,\, \bar{Q}_t^{(P)} < K\right\}\right] \\
&\quad + \mathbb{E}\left[meas\left\{t \in [0,T] : k(n) < \frac{1}{n}Q_t^{(P,n)} \leq K,\, \bar{Q}_t^{(P)} = K\right\}\right] \\
&\quad + \mathbb{E}\left[meas\left\{t \in [0,T] : k(n) < \frac{1}{n}Q_t^{(P,n)} \leq K,\, \bar{Q}_t^{(P)} > K\right\}\right].
\end{aligned} \tag{4.6.4}$$

Let us focus on each of the terms in the above sum separately.

Term 1. We claim that, in fact,

$$meas\left\{t \in [0,T] : k(n) < \frac{1}{n}Q_t^{(P,n)} \leq K, \bar{Q}_t^{(P)} < K\right\} \to 0, \text{ a.s., as } n \to \infty.$$

Let us temporarily fix an $\epsilon > 0$ and introduce $A_\epsilon = \{t \in [0,T] : \bar{Q}_t^{(P)} + 3\epsilon < K\}$. The measure in the last display can be rewritten as

$$\int_0^T \mathbf{1}_{\{k(n)<\frac{1}{n}Q_t^{(P,n)}\leq K,\, K-3\epsilon\leq \bar{Q}_t^{(P)}<K\}}dt + \int_{A_\epsilon} \mathbf{1}_{\{k(n)<\frac{1}{n}Q_t^{(P,n)}\leq K\}}dt. \tag{4.6.5}$$

Since $k(n) \to K$ as $n \to \infty$ there exists an $n_1 \in \mathbb{N}$, such that

$$k(n) > K - \epsilon, \text{ for all } n \geq n_1. \tag{4.6.6}$$

On the other hand, since $\frac{1}{n}Q^{(P,n)} \to \bar{Q}^{(P)}$, a.s., uniformly on compacts, i.e., since

$$\left\|\frac{1}{n}Q^{(P,n)} - \bar{Q}^{(P)}\right\|_T \to 0, \text{ a.s.,}$$

there exists a set $\Omega^* \in \mathcal{F}$ such that $\mathbb{P}[\Omega^*] = 1$ and

$$\left\|\frac{1}{n}Q^{(P,n)}(\omega) - \bar{Q}^{(P)}\right\|_T \to 0, \text{ for every } \omega \in \Omega^*.$$

In other words, for $\epsilon > 0$ fixed above and for every $\omega \in \Omega^*$, there exists an index $n(\omega)$ such that

$$\left\| \frac{1}{n} Q^{(P,n)}(\omega) - \bar{Q}^{(P)} \right\|_T < \epsilon, \text{ for every } n \geq n(\omega).$$

In particular, recalling the choice of n_1 from (4.6.6), we have that

$$\left\| \frac{1}{n} Q^{(P,n)}(\omega) - \bar{Q}^{(P)} \right\|_T < \epsilon, \text{ for every } n \geq n(\omega) \vee n_1.$$

Therefore, for every $n \geq n(\omega) \vee n_1$ we have

$$\left| \frac{1}{n} Q_t^{(P,n)}(\omega) - \bar{Q}_t^{(P)} \right| < \epsilon, \text{ for every } t \in [0,T].$$

Consequently, the following statement holds true for every $n \geq n(\omega) \vee n_1$:

$$\frac{1}{n} Q_t^{(P,n)}(\omega) < \epsilon + \bar{Q}_t^{(P)}, \text{ for every } t \in A_\epsilon.$$

By definition of the set A_ϵ, the above inequality implies that for every $n \geq n(\omega) \vee n_1$,

$$\frac{1}{n} Q_t^{(P,n)}(\omega) < K - 2\epsilon, \text{ for every } t \in A_\epsilon.$$

Recalling the choice of the index n_1 from (4.6.6), we get that the last inequality in turn yields that for every $n \geq n(\omega) \vee n_1$,

$$\frac{1}{n} Q_t^{(P,n)}(\omega) < k(n), \text{ for every } t \in A_\epsilon.$$

Thus, for all $\omega \in \Omega^*$ and for every $n \geq n(\omega) \vee n_1$ we have

$$\int_{A_\epsilon} \mathbf{1}_{[k(n) < \frac{1}{n} Q_t^{(P,n)}(\omega) \leq K]} \, dt = 0.$$

Therefore,

$$\int_{A_\epsilon} \mathbf{1}_{[k(n) < \frac{1}{n} Q_t^{(P,n)}(\omega) \leq K]} \, dt \to 0, \ a.s.$$

Letting ϵ go to zero will make the first integral in (4.6.5), which is dominated by

$$\int_0^T \mathbf{1}_{\{K - 3\epsilon \leq \bar{Q}_t^P < K\}} \, dt,$$

a bound independent of n, vanish as well.

Term 2. Here we aim to prove that

$$\mathbb{E}\left[meas \left\{ t \in [0,T] : k(n) < \frac{1}{n} Q_t^{(P,n)} \leq K, \, \bar{Q}_t^{(P)} = K \right\} \right] \to 0.$$

Let us define $C = \{t \in [0,T] : \bar{Q}_t^{(P)} = K\}$. Furthermore, let $B_n = (\sqrt{n}(k(n) - K), 0]$, for every n and $\hat{Q}_t^{(P,n)} = \sqrt{n}(\frac{1}{n}Q_t^{(P,n)} - K)$, for every n and t.

Using Fubini's theorem and rearranging the terms, we have

$$\mathbb{E}\left[\int_C \mathbf{1}_{\{k(n)<\frac{1}{n}Q_t^{(P,n)}\leq K\}} dt\right] = \int_C \mathbb{P}[\hat{Q}_t^{(P,n)} \in B_n]dt, \tag{4.6.7}$$

where $\hat{Q}_t^{(P,n)} = \sqrt{n}(\frac{1}{n}Q_t^{(P,n)} - K)$, for all t. Noting that

$$\sqrt{n}(K - k(n)) = \sqrt{n}\left(K - \frac{1}{n}(\lfloor nK_1\rfloor + \lfloor nK_2\rfloor)\right) \leq \sqrt{n}\left(K - \frac{1}{n}(nK_1 - 1 + nK_2 - 1)\right) = \frac{2}{\sqrt{n}},$$

we conclude that the sequence $\{B_n\}$ converges to the singleton $\{0\}$. Therefore, we have for all $t \in C$,

$$\limsup_{n\to\infty} \mathbb{P}\left[\hat{Q}_t^{(P,n)} \in B_n\right] \leq \limsup_{n\to\infty} \mathbb{P}\left[\hat{Q}_t^{(P,n)} \in B_m\right],$$

for any fixed $m \in \mathbb{N}$. According to Corollary D.3.5, we get that for any m

$$\limsup_{n\to\infty} \mathbb{P}\left[\hat{Q}_t^{(P,n)} \in B_m\right] = \mathbb{P}[Y_t \in B_m],$$

with Y_t a diffuse random variable (its exact distribution is not relevant at present). What remains to be done is to let m go to infinity, and note that $B_m \to \{0\}$. This implies that the integrand on the right-hand side of (4.6.7) converges pointwise to zero, and allows Lebesgue's dominated convergence theorem to wrap up the proof.

Term 3. This proof would be a repetition of what is done in the proof for *Term 1.*, and so let us omit it. □

4.6.2 The Lower Bound

The pooled system just discussed is a suitable tool for establishing a lower bound on the performance measures $J^{(n)}$, in a manner analogous to the fluid limit case.

Lemma 4.6.7. *For every $n \in \mathbb{N}$, all $\mu \in \mathcal{L}^{(n)}(m)$ and for every $t \in [0,T]$, $L_t^{(2,n)}(\mu) \geq L_t^{(P,n)}$, almost surely.*

Proof. For every $\mu \in \mathcal{L}^{(n)}(m)$, we have for all n and t

$$L_t^{(2,n)}(\mu) = \sup_{s\leq t}[-N_1^-(n\mathcal{I}_s(\mu)) + L_s^{(1,n)}(\mu) + N_2^-(n\mathcal{I}_s(\mu^2))], \; a.s.$$

The above equality implies that

$$\begin{aligned} L_t^{(2,n)}(\mu) &= \sup_{s\leq t}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu)) + L_s^{(1,n)}(\mu) - N_1^+(n\mathcal{I}_s(\lambda)) + N_2^-(n\mathcal{I}_s(\mu^2))] \\ &= \sup_{s\leq t}[Q_s^{(1,n)} - N_1^+(n\mathcal{I}_s(\lambda)) + N_2^-(n\mathcal{I}_s(\mu^2))]. \end{aligned}$$

almost surely, for every n and t. Since $Q^{(1,n)}$ is by definition nonnegative, the last display yields

$$L_t^{(2,n)}(\mu) \geq \sup_{s\leq t}[-N_1^+(n\mathcal{I}_s(\lambda)) + N_2^-(n\mathcal{I}_s(\mu^2))] = L_t^{(P,n)},$$

almost surely, for every n and t. □

Lemma 4.6.8. *For every $n \in \mathbb{N}$ and every $\mu \in \mathcal{L}^{(n)}(m)$, we have that*

$$Q^{(1,n)}(\mu) + Q^{(2,n)}(\mu) \geq Q^{(P,n)},\ a.s.$$

Moreover, for every random time $T' \in [0,T]$

$$\int_0^{T'} \left[\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,n)}(\mu)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,n)}(\mu)>K_2\}}\right] dt \geq meas\left\{t \in [0,T'] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\},\ a.s.$$

Proof. According to (4.4.1) and Lemma 4.6.7, for every $n \in \mathbb{N}$ and every $\mu \in \mathcal{L}^{(n)}(m)$, we have that for every $t \in [0,T]$,

$$\begin{aligned} Q_t^{(1,n)}(\mu) + Q_t^{(2,n)}(\mu) &= N_1^+(n\mathcal{I}_t(\lambda)) - N_2^-(n\mathcal{I}_t(\mu^2)) + L_t^{(2,n)}(\mu) \\ &\geq N_1^+(n\mathcal{I}_t(\lambda)) - N_2^-(n\mathcal{I}_t(\mu^2)) + L_t^{(P,n)} = Q_t^{(P,n)}. \end{aligned}$$

This is exactly the first claim announced in the lemma.

Furthermore, regardless of the choice of μ, for all t such that $\frac{1}{n}Q_t^{(P,n)} > K_1 + K_2$, necessarily either $\frac{1}{n}Q_t^{(1,n)}(\mu) > K_1$ or $\frac{1}{n}Q_t^{(2,n)}(\mu) > K_2$, yielding the announced lower bound. □

Applying Lemma 4.6.8 to $T' = T$, we obtain the final result of this subsection.

Proposition 4.6.9. *For every $n \in \mathbb{N}$ and every $\mu \in \mathcal{L}^{(n)}(m)$*

$$J^{(n)}(\mu) \geq meas\left\{t \in [0,T] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\}.$$

Let us complete the discussion by formally defining the sequence $\{J_{LB}^{(n)}\}$, representing the just proposed lower bound,

$$J_{LB}^{(n)} = meas\left\{t \in [0,T] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\}. \tag{4.6.8}$$

4.6.3 Definition of the Proposed Optimal Discipline

Let us begin by fixing a function $g : \mathbb{N} \to \mathbb{R}_+$ satisfying the following assumption.

Assumption 4.6.10. For a given fixed constant $\epsilon > 0$

(i) As $n \to \infty$, $\frac{g(n)}{n^{1+\epsilon}} \to \infty$.

(ii) As $n \to \infty$, $n^{1+\epsilon} \ln(\frac{g(n)}{g(n)+n}) \to 0$.

We restrict our attention for now to the case $m = \infty$, i.e., the case in which the amount of service at our disposal is unbounded, and define the sequence $\{\mu_n\}$ of admissible service disciplines as

$$\mu_n(t) := g(n)[(\lambda_t + \mu_t^2) \vee 1]\mathbf{1}_{\{Q_t^{(1,n)}(\mu_n)>0,\ Q_t^{(2,n)}(\mu_n)<\lfloor nK_2 \rfloor\}}, \tag{4.6.9}$$

for every $n \in \mathbb{N}$ and $t \in [0, T]$. In words, the proposed policies do not serve when the first queue is empty (i.e., they are "work-conserving") or when one more arrival into the second queue would cause the queue to cross over the threshold; otherwise, the rate of service of proposed policies is large enough to (asymptotically) cause that the services in the first queue precede the next arrival into the first station and the next departure from the second one.

In the exposition to follow (i.e., in Subsections 4.6.4-4.6.7) we focus solely on the performance of the system once the sequence of disciplines $\{\mu_n\}$ is employed. So we omit the dependence on the service discipline when referring to the queue lengths and regulator processes.

4.6.4 Bound on the Waiting Time at B_n

For every n, we denote the space of all possible pairs of coordinates representing the lengths of $Q^{(1,n)}$ and $Q^{(2,n)}$ by

$$S^{(n)} = \left\{ (q_1, q_2) : q_1 = \frac{k_1}{n},\ q_2 = \frac{k_2}{n},\ k_1, k_2 \in \mathbb{N}_0 \right\}. \tag{4.6.10}$$

Next, we partition this set as $S^{(n)} = S_w^{(n)} \cup S_r^{(n)} \cup S_p^{(n)}$, with $S_w^{(n)}, S_r^{(n)}$ and $S_p^{(n)}$ defined in the following way:

$$\begin{aligned} S_r^{(n)} &= \left\{ (q_1, q_2) \in S^{(n)} : q_1 \le \frac{1}{n}\lfloor nK_1 \rfloor, q_2 \le \frac{1}{n}\lfloor nK_2 \rfloor \right\}, \\ S_p^{(n)} &= \left\{ (q_1, q_2) \in S^{(n)} : q_1 + q_2 > \frac{1}{n}(\lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor) \right\}, \\ S_w^{(n)} &= S^{(n)} - (S_p^{(n)} \cup S_r^{(n)}). \end{aligned}$$

For every n, we denote by B_n a particular element of the boundary of $S_w^{(n)}$. Namely, we set $B_n = (\frac{1}{n}(\lfloor nK_1 \rfloor + 1), \frac{1}{n}(\lfloor nK_2 \rfloor - 1))$. We will be interested in the amount of time the system spends at B_n. So let us introduce two sequences of stopping times - $\{\sigma_i^{(n)}\}_i$ representing the times the system begins a visit to the point B_n, and $\{\zeta_i^{(n)}\}_i$ representing the moments of exit from B_n. Formally, we set

$$\begin{aligned} \sigma_1^{(n)} &= \inf\{t > 0 : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\} \wedge T, \\ \zeta_i^{(n)} &= \inf\{t > \sigma_i^{(n)} : (Q_t^{(1,n)}, Q_t^{(2,n)}) \neq nB_n\} \wedge T, \text{ for all } i \in \mathbb{N}, \\ \sigma_i^{(n)} &= \inf\{t > \zeta_{i-1}^{(n)} : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\} \wedge T, \text{ for all } i > 1. \end{aligned}$$

Now, the amount of time the system spends at B_n can be expressed as

$$meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\} = \sum_{i\in\mathbb{N}} (\zeta_i^{(n)} - \sigma_i^{(n)}).$$

By the strong Markov property of the state process $(Q^{(1,n)}, Q^{(2,n)})$, the random variables $\{\zeta_i^{(n)} - \sigma_i^{(n)}\}_i$ are mutually independent.

Let the sequence of random variables $\{\theta_i^{(n)}\}_i$ denote the service times in the first station at each visit to B_n. The terms in the sequence $\{\zeta_i^{(n)} - \sigma_i^{(n)}\}_i$ are all dominated by the terms in the sequence $\{\theta_i^{(n)}\}_i$. Furthermore, the sequence $\{\theta_i^{(n)}\}_i$ is independent of the arrival process N_1^+, by the assumptions of independence imposed on the processes N_1^+ and N_1^-.

Let us continue by partitioning the space of all possible trajectories for a given n into measurable sets $A^{(n)}$ and $(A^{(n)})^c$, with $A^{(n)}$ defined as

$$A^{(n)} = \{\omega \in \Omega : N_1^+(n\mathcal{I}_T(\lambda))(\omega) \le n^{1+\epsilon}\}, \tag{4.6.11}$$

with ϵ the positive constant fixed in Assumption 4.6.10.

Lemma 4.6.11. *For every n and i, we have that $\mathbb{E}[\theta_i^{(n)}] \le \frac{1}{ng(n)}$.*

Proof. Let us temporarily fix the index n. Then, for every i, the random variable $\theta_i^{(n)}$ is non-negative. Therefore, its expected value can be rewritten as

$$\mathbb{E}\left[\theta_i^{(n)}\right] = \int_0^\infty \mathbb{P}[\theta_i^{(n)} > t]\, dt. \tag{4.6.12}$$

Let us introduce the sequence of stopping times υ_i denoting the starting times of visits to nB_n of the state process $(Q^{(1,n)}, Q^{(2,n)})$, i.e., let

$$\begin{aligned} \upsilon_1 &= \inf\{t > 0 : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\} \wedge T, \\ \hat{\upsilon}_i &= \inf\{t > \upsilon_i : (Q_t^{(1,n)}, Q_t^{(2,n)}) \neq nB_n\} \wedge T, \\ \upsilon_i &= \inf\{t > \hat{\upsilon}_{i-1} : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\} \wedge T. \end{aligned}$$

Then, for every i and t, we have that

$$\mathbb{P}\left[\theta_i^{(n)} > t\right] = \mathbb{P}\left[N_1^-\left(ng(n)\int_{\upsilon_i}^{t+\upsilon_i} (\lambda_s + \mu_s^2) \vee 1\, ds\right) = 0\right].$$

Thanks to the monotonicity of the Poisson process N_1^-, we can bound the expression above as follows:

$$\mathbb{P}\left[\theta_i^{(n)} > t\right] \le \mathbb{P}\left[N_1^-(ng(n)t) = 0\right] = e^{-ng(n)t}.$$

Thus the expected value in (4.6.12) can be bounded from above in the following manner:

$$\mathbb{E}\left[\theta_i^{(n)}\right] \le \int_0^\infty e^{-ng(n)t}\,dt = \frac{1}{ng(n)}.$$

□

Lemma 4.6.12. *As* $n \to \infty$, $\mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{A^{(n)}}] \to 0$.

Proof. Provided that the total number of arrivals into the system does not exceed $n^{1+\epsilon}$, the total number of visits to B_n cannot exceed $n^{1+\epsilon}$. Hence, for all n,

$$\begin{aligned}
&\mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{A^{(n)}}] \\
&\quad = \mathbb{E}\left[\mathbf{1}_{A^{(n)}} \sum_{1 \le i \le n^{1+\epsilon}} (\zeta_i^{(n)} - \sigma_i^{(n)})\right] \\
&\quad \le \mathbb{E}\left[\sum_{1 \le i \le n^{1+\epsilon}} (\zeta_i^{(n)} - \sigma_i^{(n)})\right] \\
&\quad \le \mathbb{E}\left[\sum_{1 \le i \le n^{1+\epsilon}} \theta_i^{(n)}\right] \\
&\quad = \sum_{1 \le i \le n^{1+\epsilon}} \mathbb{E}[\theta_i^{(n)}].
\end{aligned}$$

Due to Lemma 4.6.11 and the last inequality, we conclude that

$$\mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{A^{(n)}}] \le \frac{n^{1+\epsilon}}{ng(n)} = \frac{n^\epsilon}{g(n)}.$$

Invoking Assumption 4.6.10, we get the desired result. □

Lemma 4.6.13. *As* $n \to \infty$, $\mathbb{P}[A^{(n)}] \to 1$.

Proof. By definition, for all n, we have

$$1 - \mathbb{P}[A^{(n)}] = \mathbb{P}[(A^{(n)})^c] = \mathbb{P}[N_1^+(n\mathcal{I}_T(\lambda))(\omega) > n^{1+\epsilon}] \le \frac{n\mathcal{I}_T(\lambda)}{n^{1+\epsilon}} = \frac{\mathcal{I}_T(\lambda)}{n^\epsilon}.$$

Letting $n \to \infty$ completes the argument. □

Proposition 4.6.14. *As* $n \to \infty$,

$$\mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}] \longrightarrow 0.$$

Proof. For all n,

$$\begin{aligned}\mathbb{E}[meas\{t \in [0,T] :&(Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}] \\ &= \mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{A^{(n)}}] \\ &\quad + \mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{(A^{(n)})^c}] \\ &\leq \mathbb{E}[meas\{t \in [0,T] : (Q_t^{(1,n)}, Q_t^{(2,n)}) = nB_n\}\mathbf{1}_{A^{(n)}}] + T\mathbb{P}[(A^{(n)})^c].\end{aligned}$$

By Lemmas 4.6.12 and 4.6.13, the sum on the right-hand side vanishes as $n \to \infty$. □

4.6.5 The Sequence of Transitions

Let us, for all n, define the sequence $\{\eta_i^{(n)}\}_i$ of stopping times representing the instances at which the system driven by service disciplines $\{\mu_n\}$ changes position; namely, let

$$\begin{aligned}\eta_1^{(n)} &= \inf\{t > 0 : Q_t^{(1,n)} > 0\} \wedge T, \\ \eta_i^{(n)} &= \inf\{t > \eta_{i-1}^{(n)} : (Q_t^{(1,n)}, Q_t^{(2,n)}) \neq (Q^{(1,n)}{}_{\eta_{i-1}^{(n)}}, Q^{(2,n)}{}_{\eta_{i-1}^{(n)}})\} \wedge T, \text{ for } i > 1.\end{aligned}$$

Among the transitions of the system, we wish to emphasize the difference between the ones created by a service in the first queue, and the ones occurring due to either an arrival into the first queue or a completion of service in the second queue. Hence, let us define a sequence of random variables $\{Z_i^{(n)}\}_i$, for every n, by

$$\begin{aligned}Z_1^{(n)} &= 0, \\ Z_i^{(n)} &= \begin{cases} 0 & \text{if } Q^{(1,n)}{}_{\eta_{i-1}^{(n)}} \leq Q^{(1,n)}{}_{\eta_i^{(n)}} \text{ or } Q^{(2,n)}{}_{\eta_{i-1}^{(n)}} \geq Q^{(2,n)}{}_{\eta_i^{(n)}} \\ 1 & \text{otherwise} \end{cases} \quad \text{for all } i > 1.\end{aligned}$$

For every n, we focus on indices i not exceeding $2n^{1+\epsilon}$ such that there is non-zero service in the first queue at the time preceding the i^{th} transition in the system. More formally, let us define a sequence of sets $\{C^{(n)}\}$ as

$$C^{(n)} = \{i \leq 2n^{1+\epsilon} : Q^{(1,n)}{}_{\eta_{i-1}^{(n)}} > 0 \text{ and } Q^{(2,n)}{}_{\eta_{i-1}^{(n)}} < \lfloor nK_2 \rfloor\}.$$

We are interested in a bound on the probability of the existence of an instance among the first $2n^{1+\epsilon}$ transitions preceded by a period of non-zero service in the first queue, such that the transition i does **not** occur due to the completion of service in the first queue.

Lemma 4.6.15. *Suppose Assumption 4.6.10 holds. As $n \to \infty$, $\mathbb{P}[\exists i \in C^{(n)}, Z_i^{(n)} = 0] \to 0$.*

Proof. For all n,

$$\mathbb{P}[\exists i \in C^{(n)}, Z_i^{(n)} = 0] = 1 - \mathbb{P}[\forall i \in C^{(n)}, Z_i^{(n)} = 1].$$

By the strong Markov property of the state process $(Q^{(1,n)}, Q^{(2,n)})$, the random variables $\{Z_i^{(n)}\}_{i \in C^{(n)}}$ are independent.

Furthermore, for all $i \in C^{(n)}$, the event $\{Z_i^{(n)} = 1\}$ can be described as the event in which the random variable governing the service times in the first queue occurs before an arrival into the first queue or a completion of service in the second one. Let us note that the latter two events of interest can be regarded as emanating from a single unit Poisson process independent from N_1^-, and with rate $\lambda + \mu^2 \in \mathbb{L}_+^1[0,T]$.

Invoking the strong Markov Property of the state process $(Q^{(1,n)}, Q^{(2,n)})$ once again, we realize that, conditionally on $C^{(n)}$, we can interpret

(i) the service process in the first station and

(ii) the Poisson process with rate $\lambda + \mu^2$ which counts both the arrivals into the first station and the departures from the second station

as at most $n^{1+\epsilon}$ independent pairs of Poisson processes we consider in Corollary D.1.3. Taking into account Assumption 4.6.10, we see that all the conditions of Corollary D.1.3 are satisfied, which completes the proof of this lemma. □

4.6.6 Bound on the Time Spent in $S_w^{(n)}$

Let us start by noting that it is straightforward from the definition of $\{\mu_n\}$ given in (4.6.9) that the system is never going to reach the set $S_w^{(n)} \cap \{(q_1, q_2) \in S^{(n)} : q_2 > \frac{1}{n}\lfloor nK_2 \rfloor\}$.

Lemma 4.6.16. *As $n \to \infty$,*

$$\mathbb{P}\left[\left\{\exists t \in [0,T] \text{ such that } \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\} \cap A^{(n)}\right] \longrightarrow 0.$$

Proof. Looking at the possible transitions in the grid $S^{(n)}$, it is evident that a necessary condition for the system to ever visit $S_w^{(n)} \setminus \{B_n\}$ is that a service in the first station took a longer time to complete then either a new arrival into the first station or a service in the second station. This event is necessary in order for second coordinate of the state process $(Q^{(1,n)}(\mu_n), Q^{(2,n)}(\mu_n))$ to be at $\lfloor nK_2 \rfloor - 1$ while the first coordinate is over its threshold, i.e., at $\lfloor nK_1 \rfloor + 1$. Formally, this is the event that the random variable $Z_i^{(n)}$ assumes a value of 0 for some $i \in C^{(n)}$. Therefore,

$$\mathbb{P}\left[\left\{\exists t, \text{ such that } \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\} \cap A^{(n)}\right] \leq \mathbb{P}[\exists i \in C^{(n)}, Z_i^{(n)} = 0].$$

Using Lemma 4.6.15 completes the proof. □

We are now ready to prove that the expected amount of time the system spends in the set $S_w^{(n)} \setminus \{B_n\}$ vanishes altogether.

Proposition 4.6.17. *As* $n \to \infty$*,*

$$\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\}\right] \to 0.$$

Proof. The expression in the proposition equals the sum

$$\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\} \mathbf{1}_{A^{(n)}}\right]$$
$$+ \mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\} \mathbf{1}_{(A^{(n)})^c}\right],$$

for all n. The first term in the sum is dominated by

$$T\,\mathbb{P}\left[\left\{\exists t,\ \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)} \setminus \{B_n\}\right\} \cap A^{(n)}\right],$$

and it disappears as $n \to \infty$, due to Lemma 4.6.16. The second term is, on the other hand, dominated by $T\mathbb{P}[(A^{(n)})^c]$, which also vanishes in the limit, as is seen in Lemma 4.6.13. □

When the last proposition is combined with Lemma 4.6.12, we obtain the following result.

Corollary 4.6.18. *As* $n \to \infty$*,*

$$\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q_t^{(1,n)}, Q_t^{(2,n)}) \in S_w^{(n)}\right\}\right] \to 0.$$

4.6.7 The Asymptotic Lower Bound on the Grids

Lemma 4.6.19. *For every* $n \in \mathbb{N}$ *and every* $\mu \in \mathcal{L}^{(n)}(m)$*, we have that*

$$J^{(n)}(\mu) \geq meas\{t \in [0,T] : Q_t^{(P,n)} > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor\},\ a.s.$$

Proof. It suffices to prove that for all n, all μ and all $t \in [0,T]$, we have

$$Q_t^{(P,n)} > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor \implies Q_t^{(1,n)}(\mu) > \lfloor nK_1 \rfloor \text{ or } Q_t^{(2,n)}(\mu) > \lfloor nK_2 \rfloor,$$

almost surely. This implication, however follows straightforwardly from the fact that, by Lemma 4.6.8,

$$Q_t^{(P,n)} \leq Q_t^{(1,n)}(\mu) + Q_t^{(2,n)}(\mu),$$

across all choices for n, μ and t. □

Let us introduce, for all n, the random variable $J^{(n)}_{DLB}$ denoting the lower bound obtained in the last lemma, i.e., let

$$J^{(n)}_{DLB} = meas\{t \in [0,T] : Q^{(P,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor\}. \tag{4.6.13}$$

Next, we want to verify that the expected values of the sequence of bounds $\{J^{(n)}_{DLB}\}$ involving the pooled queue matches the expected amount of time the pair of queue lengths in the tandem spends in the region $nS^{(n)}_p$, provided that the sequence of service disciplines $\{\mu_n\}$ from the defining expression (4.6.9) is used.

Lemma 4.6.20. *Let $\{\mu_n\}$ be the sequence of service disciplines defined in (4.6.9). Then, as $n \to \infty$,*

$$\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q^{(1,n)}_t, Q^{(2,n)}_t) \in S^{(n)}_p\right\} - J^{(n)}_{DLB}\right] \to 0. \tag{4.6.14}$$

Proof. For every n, the expectation in (4.6.14) can be expanded in the following way, using the definitions of the set $S^{(n)}_p$ and the random variable $J^{(n)}_{DLB}$,

$$\begin{aligned}
&\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q^{(1,n)}_t, Q^{(2,n)}_t) \in S^{(n)}_p\right\} - J^{(n)}_{DLB}\right] \\
&\quad = \mathbb{E}[meas\{t \in [0,T] : Q^{(1,n)}_t + Q^{(2,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor\}] \\
&\qquad - \mathbb{E}[meas\{t \in [0,T] : Q^{(P,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor\}].
\end{aligned}$$

Applying the Fubini-Tonelli theorem to both integrals, as both integrands are nonnegative we get that the above expectation equals

$$\int_0^T (\mathbb{P}[Q^{(1,n)}_t + Q^{(2,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor] - \mathbb{P}[Q^{(P,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor])\, dt.$$

Thus, due to the first result of Lemma 4.6.8 and the above calculation, we have that the expected value in (4.6.14) is equal to

$$\int_0^T \mathbb{P}[Q^{(1,n)}_t + Q^{(2,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor, Q^{(P,n)}_t \leq \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor]\, dt. \tag{4.6.15}$$

For every t, the integrand in (4.6.15) can be dealt with as follows

$$\mathbb{P}\left[Q^{(1,n)}_t + Q^{(2,n)}_t > \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor, Q^{(P,n)}_t \leq \lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor\right] \leq \mathbb{P}\big[Q^{(1,n)}_t + Q^{(2,n)}_t \neq Q^{(P,n)}_t\big].$$

The right-hand side of the last display converges to zero as $n \to \infty$ by Lemma D.4.1. Using Lebesgue's Dominated Convergence Theorem, we conclude that the resulting integral in (4.6.15) vanishes as well when $n \to \infty$. This wraps up the proof. □

With the last lemma completed, we are ready to prove that the difference of the expected performance of $\{\mu_n\}$ and the expected values of the lower bounds $J^{(n)}_{DLB}$ disappears in the limit.

Proposition 4.6.21. *As $n \to \infty$,*

$$\mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}_{DLB}] \to 0,$$

where the sequence $\{\mu_n\}$ is as defined in (4.6.9).

Proof. Recalling the definition of $S^{(n)}_w$, for all n,

$$\begin{aligned}\mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}_{DLB}] = \mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q^{(1,n)}_t, Q^{(2,n)}_t) \in S^{(n)}_w\right\}\right] \\ + \mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(Q^{(1,n)}_t, Q^{(2,n)}_t) \in S^{(n)}_p\right\} - J^{(n)}_{DLB}\right].\end{aligned}$$

The claims of Corollary 4.6.18 and Lemma 4.6.20 together complete the proof. ☐

We will next establish that the expected value of the difference between the lower bounds $J^{(n)}_{LB}$ and their discrete analogues $J^{(n)}_{DLB}$ vanishes as $n \to \infty$. This will allow us to focus on the grid $S^{(n)}$, for any n.

Lemma 4.6.22. *Let $J^{(n)}_{LB}$ and $J^{(n)}_{DLB}$ be defined as in (4.6.8) and (4.6.13), respectively. Then, as $n \to \infty$, $\mathbb{E}[J^{(n)}_{LB} - J^{(n)}_{DLB}] \to 0$.*

Proof. Directly from the definitions of $J^{(n)}_{LB}$ and $J^{(n)}_{DLB}$, for all n the expectation of their difference equals

$$\mathbb{E}\left[meas\left\{t \in [0,T] : \frac{1}{n}(\lfloor nK_1 \rfloor + \lfloor nK_2 \rfloor) < \frac{1}{n}Q^{(P,n)}_t \leq K_1 + K_2\right\}\right].$$

The announced convergence is precisely the content of Lemma 4.6.6. ☐

Theorem 4.6.23. *As $n \to \infty$,*

$$\mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}_{LB}] \longrightarrow 0.$$

Proof. For all n, we can rewrite the above expectation as

$$\mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}_{LB}] = \mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}_{DLB}] - \mathbb{E}[J^{(n)}_{LB} - J^{(n)}_{DLB}].$$

Combining Proposition 4.6.21 and Lemma 4.6.22, we obtain the result. ☐

4.6.8 The nontrivial constraint $m < \infty$

Recall that the constraint imposed on the admissible policies is $I_T(\mu) \le m$ for some given constant $m \in \mathbb{R} \cup \infty$. Let us begin with the definition of the sequence of random times $\{\tau^{(n)}\}$ given by

$$\tau^{(n)} = \inf\{t \in [0,T] : N_1^+(n\mathcal{I}_t(\lambda)) > N_1^-(nm) + nK_1\} \wedge T. \tag{4.6.16}$$

These times play a role analogous to the one of time $\tau(K_1+m)$ in the fluid limit analysis - more precisely, in Lemma 4.4.3 and Corollary 4.4.4.

By Lemma 4.6.8, for every index n and every admissible control $\mu \in \mathcal{L}^{(n)}(m)$, the inequality

$$\int_0^{\tau^{(n)}} \left[\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,n)}(\mu)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,n)}(\mu)>K_2\}}\right] dt \ge meas\left\{t \in [0,\tau^{(n)}] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\} \tag{4.6.17}$$

holds almost surely. On the other hand, for every $t > \tau^{(n)}$, we have that

$$\begin{aligned} Q_t^{(1,n)}(\mu) \ge X_t^{(1,n)}(\mu) &= N_1^+(n\mathcal{I}_t(\lambda)) - N_1^-(n\mathcal{I}_t(\mu)) \\ &> N_1^-(nm) + nK_1 - N_1^-(n\mathcal{I}_t(\mu)) \ge nK_1,\ a.s. \end{aligned} \tag{4.6.18}$$

Combining inequalities (4.6.17) and (4.6.18) we obtain the following lower bound on the performance measure $J^{(n)}$ for any $\mu \in \mathcal{L}^{(n)}(m)$:

$$J^{(n)}(\mu) \ge meas\left\{t \in [0,\tau^{(n)}] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\} + T - \tau^{(n)} =: J_{LBF}^{(n)},\ a.s. \tag{4.6.19}$$

Let us recall the sequence of service disciplines given in (4.6.9) and define a sequence of stopping times as

$$\tau_*^{(n)} = \inf\{t \in [0,T] : \mathcal{I}_t(\mu_n) = m\}.$$

Next, we introduce the following sequence of admissible service disciplines (i.e., service disciplines conforming to the constraint on the available amount of service)

$$\mu_n^* = \mu_n \mathbf{1}_{[0,\tau_*^{(n)}]}. \tag{4.6.20}$$

Due to the discussion leading to inequality (4.6.19), we have that for every n

$$\begin{aligned} J^{(n)}(\mu_n^*) - J_{LBF}^{(n)} = &\int_0^{\tau^{(n)}} \mathbf{1}_{\{\frac{1}{n}Q_t^{(1,n)}(\mu_n^*)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,n)}(\mu_n^*)>K_2\}}\, dt \\ &- \text{meas}\left\{t \in [0,\tau^{(n)}] : \frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right\}. \end{aligned}$$

Using the definition of random times $\tau_*^{(n)}$, we can rewrite the above equality as

$$\begin{aligned} J^{(n)}(\mu_n^*) - J_{LBF}^{(n)} = \int_0^{\tau^{(n)}\wedge\tau_*^{(n)}} & \left(\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,\,n)}(\mu_n^*)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,\,n)}(\mu_n^*)>K_2\}}\right) dt \\ & - \text{meas}\left\{t \in [0, \tau^{(n)}\wedge\tau_*^{(n)}] : \frac{1}{n}Q_t^{(P,\,n)} > K_1 + K_2\right\} \\ & + \int_{\tau^{(n)}\wedge\tau_*^{(n)}}^{\tau^{(n)}} \left(\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,\,n)}(\mu_n^*)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,\,n)}(\mu_n^*)>K_2\}}\right) dt \\ & - \text{meas}\left\{t \in (\tau^{(n)}\wedge\tau_*^{(n)}, \tau^{(n)}] : \frac{1}{n}Q_t^{(P,\,n)} > K_1 + K_2\right\}. \end{aligned} \tag{4.6.21}$$

On the segment $[0, \tau^{(n)}\wedge\tau_*^{(n)}]$, we have that $\mu_n^* = \mu_n$. Hence, we can - *mutatis mutandis* - reemploy the argument from the present section leading to Theorem 4.6.23 to arrive at the following conclusion

$$\begin{aligned} \mathbb{E}\Big[\int_0^{\tau^{(n)}\wedge\tau_*^{(n)}} & \left(\mathbf{1}_{\{\frac{1}{n}Q_t^{(1,\,n)}(\mu_n^*)>K_1\}} + \mathbf{1}_{\{\frac{1}{n}Q_t^{(2,\,n)}(\mu_n^*)>K_2\}}\right) dt \\ & - \mathit{meas}\left\{t \in [0, \tau^{(n)}\wedge\tau_*^{(n)}] : \frac{1}{n}Q_t^{(P,\,n)} > K_1 + K_2\right\}\Big] \to 0 \text{ as } n \to \infty. \end{aligned} \tag{4.6.22}$$

As for the remaining terms on the right-hand side of (4.6.21), they do not vanish only in the case that $\tau^{(n)} > \tau_*^{(n)}$. However, in that case we have that for every $t \in [\tau_*^{(n)}, \tau^{(n)}]$

$$Q_t^{(1,\,n)}(\mu_n^*) = N_1^+(n\mathcal{I}_t(\lambda)) - N_1^-(n\mathcal{I}_t(\mu_n^*)) = N_1^+(n\mathcal{I}_t(\lambda)) - N_1^-(nm).$$

From the definition of $\tau^{(n)}$ (see (4.6.16)), we conclude that $Q^{(1,\,n)}(\mu_n^*) \le K_1$ on $[\tau_*^{(n)}, \tau^{(n)}]$. It is clear from the construction of μ_n^* that $Q^{(2,\,n)} \le K_2$ at all times. Therefore, there is no penalty accumulated over the interval $[\tau_*^{(n)}, \tau^{(n)}]$.

On the other hand, the pooled queue is bounded from above by the sum of the two queue lengths in the tandem system. Thus,

$$\text{meas}\left\{t \in (\tau^{(n)}\wedge\tau_*^{(n)}, \tau^{(n)}] : \frac{1}{n}Q_t^{(P,\,n)} > K_1 + K_2\right\} = 0.$$

To sum up, there is no contribution to either the penalty or the lower bound over the interval $[\tau_*^{(n)}, \tau^{(n)}]$. Thanks to this fact, as well as the expansion in (4.6.21) and the limit in (4.6.22), we conclude that

$$\mathbb{E}\left[J^{(n)}(\mu_n^*) - J_{LBF}^{(n)}\right] \to 0 \text{ as } n \to \infty.$$

Due to Proposition 4.6.3 applied to the last display along with (4.6.18) and (4.6.19), we can infer that the sequence $\{\mu_n^*\}$ is indeed asymptotically optimal. One simulated trajectory of the state process $(Q^{(1,n)}, Q^{(2,n)})$ is displayed in Figure 4.2. One should note that at the point when the

Figure 4.2: The Tandem System: Asymptotic optimality in the case of infinite buffers

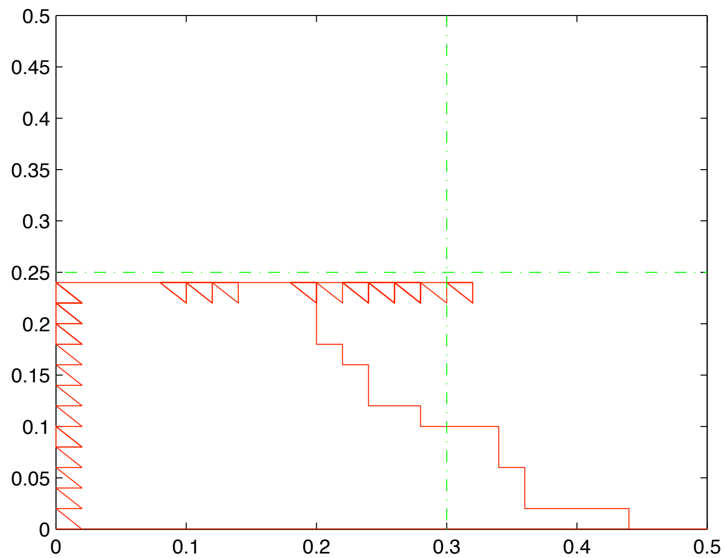

constraint on the available service is reached the trajectory starts moving in the "south-east" direction. Before that time, the structure of the trajectory is "triangular" which illustrates the fact that whenever the state process is strictly within the rectangle $[0, nK_1] \times [0, nK_2]$ and while the constraint on the total amount of service is not reached, a completion of service in the first station happens before both a new arrival in the first station and a completion of service in the second station.

Remark 4.6.1. The time $\tau^{(n)}$ is not chosen *ad hoc*. As n increases, the service disciplines $\{\mu_n^*\}$ approach optimal performance. Among other things, this means that the amount of upward pushing in the first station vanishes in the limit, i.e., serving "in vain" disappears. Thus, asymptotically the first queue will not cross over the threshold between the moment it reaches the total service constraint and the time $\tau^{(n)}$ (if $\tau^{(n)}$ is a later time).

4.7 Stochastic vs. Deterministic Service Disciplines

The previous section brought forth a description of a class of asymptotically optimal sequences of **stochastic** service disciplines. Obviously, a family of **deterministic** asymptotically optimal sequences would be preferable, as the controller could then propose a service up-front without taking the state of the system throughout the $[0, T]$ cycle into consideration at all. We prove this is impossible for certain values of parameters m, λ and μ^2.

For the sake of simplicity we immediately assume $m = \infty$, i.e., there is no constraint on the total amount of service at our disposal. Recalling the notion of asymptotic optimality from Definition 4.6.1, as well as the results leading to Theorem 4.6.23, we see that it suffices to find values of λ and μ^2 and a positive constant π, such that for all deterministic admissible sequences $\{\mu_n\}$

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J_{LB}^{(n)}] \geq \pi. \tag{4.7.1}$$

Next, we suppose the system parameters satisfy the following condition .

Assumption 4.7.1. $\mathcal{I}(\lambda) \geq \mathcal{I}(\mu^2)$

Taking a closer look at the sequence of lower bounds $\{J_{LB}^{(n)}\}$, we realize that Assumption 4.7.1 simplifies the form of the lower bounds significantly, as all reflection eventually disappears in the pooled system.

Lemma 4.7.2. *Let λ and μ^2 satisfy Assumption 4.7.1. Then we have that*

$$\mathbb{E}\left[J_{LB}^{(n)} - meas\left\{t \in [0,T] : \frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right\}\right] \to 0, \text{ as } n \to \infty.$$

Proof. Using Fubini's theorem to exchange the order of the expectation and taking the Lebesgue measure in the above expression, we see that the claim of the lemma is equivalent to

$$\int_0^T \left(\mathbb{P}\left[\frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right] - \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right]\right) dt \to 0, \text{ as } n \to \infty. \tag{4.7.2}$$

By the Functional Strong Law of Large Numbers, we have that

$$\frac{1}{n}X_t^{(P,n)} \to \bar{X}^P = \mathcal{I}(\lambda - \mu^2), \text{ a.s., as } n \to \infty, \tag{4.7.3}$$

in the uniform topology. The claim in the last display, along with the continuity of the Skorokhod mapping yields

$$\frac{1}{n}Q_t^{(P,n)} \to \Gamma(\bar{X}^P) = \mathcal{I}(\lambda - \mu^2), \text{ a.s., as } n \to \infty, \tag{4.7.4}$$

uniformly. Portmanteau's Theorem applied to the results (4.7.3) and (4.7.4) gives us that

$$\lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right] = \lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right], \tag{4.7.5}$$

for all t such that $\bar{X}_t^P \neq K_1 + K_2$.

On the other hand, by Theorems 9.6.1. and 9.6.2. of [Whi02a], we have that as $n \to \infty$

$$\begin{aligned} \sqrt{n}\left(\frac{1}{n}X^{(P,n)} - \bar{X}^P\right) &\Rightarrow W(\mathcal{I}(\lambda + \mu^2)), \\ \sqrt{n}\left(\frac{1}{n}Q^{(P,n)} - \bar{X}^P\right) &\Rightarrow W(\mathcal{I}(\lambda + \mu^2)), \end{aligned} \tag{4.7.6}$$

uniformly, where the second relation uses the fact that $\bar{Q}^P = \bar{X}^P$ and W denotes the standard Brownian motion. In particular, for all t such that $\bar{X}^P_t = K_1 + K_2$, employing Portmanteau's Theorem once more (this time to obtain convergence of probabilities from convergence in distribution to a random variable that admits a density), we get that

$$\begin{aligned}\lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n}Q_t^{(P,n)} > K_1 + K_2\right] &= \lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] \\ &= \lim_{n\to\infty} \mathbb{P}[W(\mathcal{I}_t(\lambda+\mu^2)) > 0] = \frac{1}{2}.\end{aligned} \tag{4.7.7}$$

Combining the limits in (4.7.5) and (4.7.7), and applying the Lebesgue's Dominated Convergence Theorem to the integral in (4.7.2), we obtain the announced claim. □

In view of the previous lemma, we will from now redefine the lower bound to the performance measure in question to be

$$\breve{J}_{LB}^{(n)} = \text{meas}\left\{t \in [0,T] : \frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right\}. \tag{4.7.8}$$

Its expectation is given by

$$\begin{aligned}\mathbb{E}[\breve{J}_{LB}^{(n)}] &= \mathbb{E}\left[\text{meas}\left\{t \in [0,T] : \frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right\}\right] \\ &= \int_0^T \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt.\end{aligned} \tag{4.7.9}$$

We used Fubini's theorem to exchange the order of integration. The following partition of $[0,T]$ will be convenient in calculations ahead. Let

$$A = \{t \in [0,T] : \bar{X}_t^P = K_1 + K_2\}, \quad B = \{t \in [0,T] : \bar{X}_t^P < K_1 + K_2\}, \tag{4.7.10}$$

and $C = [0,T] \setminus (A \cup B)$. Equation (4.7.9) can now be rewritten as

$$\begin{aligned}\mathbb{E}[\breve{J}_{LB}^{(n)}] = &\int_A \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt + \int_B \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt \\ &+ \int_C \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt.\end{aligned} \tag{4.7.11}$$

The conclusions

$$\int_B \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt \to 0 \tag{4.7.12}$$

and

$$\int_C \mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] dt \to \text{meas}\,(C) \tag{4.7.13}$$

follow straightforwardly from the Strong Law of Large Numbers. The remaining term in (4.7.11) requires the use of the Central Limit Theorem. We begin by rewriting the integrand as follows

$$\mathbb{P}\left[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2\right] = \mathbb{P}\left[\sqrt{n}\left(\frac{1}{n}X_t^{(P,n)} - (K_1 + K_2)\right) > 0\right],$$

for all $t \in A$. By the Central Limit Theorem, $\sqrt{n}(\frac{1}{n}X_t^{(P,n)} - K_1 + K_2) \Rightarrow Y_t^P$, as $n \to \infty$, where Y_t^P is normally distributed with mean zero. Thus,

$$\mathbb{P}\left[\sqrt{n}\left(\frac{1}{n}X_t^{(P,n)} - K_1 + K_2\right) > 0\right] \to \frac{1}{2}, \text{ as } n \to \infty. \tag{4.7.14}$$

Combining (4.7.12), (4.7.13) and (4.7.14) with (4.7.11), we get

$$\mathbb{E}[\breve{J}_{LB}^{(n)}] \to \frac{1}{2}\text{meas}(A) + \text{meas}(C), \text{ as } n \to \infty. \tag{4.7.15}$$

Returning our attention to (4.7.1), we see that under Assumption 4.7.1 it suffices to provide $\lambda, \mu^2 \in \mathbb{L}_+^1[0,T]$ and $\pi > 0$, such that for all deterministic admissible $\{\mu_n\}$

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n)] \geq \pi + \frac{1}{2}\text{meas}(A) + \text{meas}(C) =: \pi'. \tag{4.7.16}$$

Let us proceed with an analysis of the inequality (4.7.16). Since the reflected queue always dominates its netput process, recalling the definition of $J^{(n)}$ from (4.3.3), a further sufficient condition for (4.7.16) is

$$\liminf_{n\to\infty} \mathbb{E}\left[\int_0^T \left(\mathbf{1}_{[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1]} + \mathbf{1}_{[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2]}\right) dt\right] \geq \pi'. \tag{4.7.17}$$

Exchanging the order of integration, we see that a sufficient condition for (4.7.17) is

$$\liminf_{n\to\infty} \int_0^T \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \geq \pi'. \tag{4.7.18}$$

Using again the partition of $[0,T]$ from (4.7.11), we can rewrite the integral on the right-hand side of (4.7.18) as

$$\begin{aligned}
&\int_0^T \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\
&\quad= \int_A \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\
&\qquad+ \int_B \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\
&\qquad+ \int_C \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt,
\end{aligned} \tag{4.7.19}$$

for any n. A trivial lower bound on the left-hand side of (4.7.19) is obtained by neglecting the integral over the domain B, which leaves us with the inequality

$$\begin{aligned}\int_0^T &\left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\ &\geq \int_A \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\ &\quad + \int_C \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt.\end{aligned} \tag{4.7.20}$$

Moreover, Fatou's Lemma applied to (4.7.20) gives us

$$\begin{aligned}\liminf_{n\to\infty} \int_0^T &\left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\ &\geq \int_A \liminf_{n\to\infty} \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\ &\quad + \int_C \liminf_{n\to\infty} \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt.\end{aligned}$$

With this, a sufficient condition for (4.7.18) is

$$\begin{aligned}&\int_A \liminf_{n\to\infty} \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \\ &\quad + \int_C \liminf_{n\to\infty} \left(\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right]\right) dt \geq \pi'.\end{aligned} \tag{4.7.21}$$

Let us introduce the following shorthand notation

$$p_t^n(\mu_n) = \mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right],$$

and expand the current notation to include the scaled centered versions of the netput processes in the prelimit sequence of tandem systems

$$\begin{aligned}\bar{X}^{(1,n)}(\mu) &= \frac{1}{n}X^{(1,n)}(\mu) - \bar{X}^{(1)}(\mu), \\ \bar{X}^{(2,n)}(\mu) &= \frac{1}{n}X^{(2,n)}(\mu) - \bar{X}^{(2)}(\mu),\end{aligned} \tag{4.7.22}$$

for all admissible μ. We next dedicate our attention to the second integrand in (4.7.21). Let t be an arbitrary instant in C. Then, by definition of the set C, we have

$$\bar{X}_t^P > K_1 + K_2. \tag{4.7.23}$$

The notation from (4.7.22) allows us to write

$$p_t^n(\mu_n) = \mathbb{P}[\bar{X}_t^{(1,n)}(\mu_n) > K_1 - \bar{X}_t^{(1)}(\mu_n)] + \mathbb{P}[\bar{X}_t^{(2,n)}(\mu_n) > K_2 - \bar{X}_t^{(2)}(\mu_n)].$$

It is of interest to find a lower bound for $\liminf_{n\to\infty} p_t^n(\mu_n)$, independent of the choice of $\{\mu_n\}$. Let us temporarily fix an arbitrary deterministic sequence $\{\mu_n\}$ and consider any convergent subsequence $\{p_t^{n_k}(\mu_{n_k})\}$ of $\{p_t^n(\mu_n)\}$, renaming this subsequence $\{q^k\}$ and calling its limit q. Setting $\epsilon_t = \frac{1}{2}(\bar{X}_t^P - (K_1 + K_2)) > 0$, and recalling the trivial equality

$$\bar{X}_t^{(1)}(\mu_n) + \bar{X}_t^{(2)}(\mu_n) = \bar{X}_t^P, \text{ for all } n,$$

we conclude that for any n, $\bar{X}_t^{(1)}(\mu_n) \geq K_1 + \epsilon_t$ or $\bar{X}_t^{(2)}(\mu_n) \geq K_2 + \epsilon_t$. Let $\{q^{k_l}\}$ be a further subsequence of $\{q^k\}$ such that one of these inequalities holds for all elements, say,

$$\bar{X}_t^{(1)}(\mu_{n_{k_l}}) \geq K_1 + \epsilon_t, \text{ for every } l. \tag{4.7.24}$$

As we will see later on, the choice of the first queue is done without loss of generality. Since $\{q^k\}$ is a convergent sequence, its subsequence $\{q^{k_l}\}$ must also be convergent with the same limit q. Considering the following consequence of (4.7.24)

$$q^{k_l} \geq \mathbb{P}[\bar{X}_t^{(1,n_{k_l})}(\mu_{n_{k_l}}) > K_1 - \bar{X}_t^{(1)}(\mu_{n_{k_l}})] \geq \mathbb{P}[\bar{X}_t^{(1,n_{k_l})}(\mu_{n_{k_l}}) > -\epsilon_t],$$

we conclude that

$$q \geq \limsup_{l\to\infty} \mathbb{P}[\bar{X}_t^{(1,n_{k_l})}(\mu_{n_{k_l}}) > -\epsilon_t].$$

However, by the Strong Law of Large Numbers,

$$\bar{X}_t^{(1,n)}(\mu_n) \to 0, \quad \text{and} \quad \bar{X}_t^{(2,n)}(\mu_n) \to 0, \tag{4.7.25}$$

almost surely. Hence, $q \geq 1$.

Having established (4.7.25), we see that the subsequence $\{q^{k_l}\}_l$ was indeed constructed without any loss of generality. The same analysis can be conducted for any convergent subsequence of $\{p_t^n(\mu_n)\}_n$, for any t and any $\{\mu_n\}$. Therefore,

$$\int_C \liminf_{n\to\infty} \left(\mathbb{P}\left[\frac{1}{n} X_t^{(1,\,n)}(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n} X_t^{(2,\,n)}(\mu_n) > K_2\right] \right) dt \geq \text{meas}(C)$$

and the still simpler sufficient condition for deterministic disciplines to perform strictly worse than the asymptotically optimal stochastic sequence is

$$\int_A \liminf_{n\to\infty} p_t^n(\mu_n)\, dt \geq \pi + \frac{1}{2}\text{meas}(A). \tag{4.7.26}$$

Focusing on the remaining domain of integration A, let us fix an arbitrary deterministic sequence $\{\mu_n\}$ and an instant $t \in A$. Again, we consider any convergent subsequence $\{p_t^{n_k}(\mu_{n_k})\}_k$ of $\{p_t^n(\mu_n)\}$. Suppressing unnecessary notation, we rename this sequence $\{q^k\}$ and call its limit q.

To ease further exposition, let us denote by α_n the values $\mathcal{I}_t(\mu_n)$, for all n. The subsequence $\{\alpha_{n_k}\}_k$ necessarily has a further subsequence $\{\alpha_{n_{k_l}}\}_l$ converging in $\bar{\mathbb{R}}_+ = \mathbb{R}_+ \cup \{\infty\}$ to some limit we call α. The sequence $\{q^{k_l}\}$, of course, must converge to q. Different cases, depending on the value of the limit α, are analyzed next.

Case 1. Let $\alpha = \infty$, i.e., let $\{\alpha_{n_{k_l}}\}_l$ be a divergent sequence. Then, because $\bar{X}^{(2)}(\mu_{n_{k_l}}) = \mathcal{I}(\mu_{n_{k_l}}) - \mathcal{I}(\mu^2)$ we have that

$$q^{k_l} \geq \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > K_2 - \alpha_{n_{k_l}} + \mathcal{I}_t(\mu^2)].$$

Since $t \in A$, it satisfies the equality $\bar{X}_t^P = \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu^2) = K_1 + K_2$. Hence, the above inequality can be rewritten as

$$\begin{aligned} q^{k_l} &\geq \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > K_2 - \alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - (K_1 + K_2)] \\ &= \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > -\alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - K_1]. \end{aligned}$$

Thanks to the limit in (4.7.25) and the fact that $\{\alpha_{n_{k_l}}\}_l$ is assumed to be a divergent sequence in this case, we conclude that the right hand side of the last display converges to 1. Therefore, $q \geq 1$.

Case 2. Let $\alpha \in \mathbb{R}$ be such that $\alpha > \mathcal{I}_t(\lambda) - K_1$ and define $\upsilon = \frac{1}{2}(\alpha - \mathcal{I}_t(\lambda) + K_1) > 0$. Similarly as in the previous case, we have

$$\begin{aligned} q^{k_l} &\geq \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > K_2 - \alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - (K_1 + K_2)]. \\ &= \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > -\alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - K_1]. \end{aligned} \tag{4.7.27}$$

For large enough l, $\alpha_{n_{k_l}} > \mathcal{I}_t(\lambda) - K_1 + \upsilon$ and, therefore, $-\upsilon > -\alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - K_1$. Thus,

$$\mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > -\alpha_{n_{k_l}} + \mathcal{I}_t(\lambda) - K_1] \geq \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > -\upsilon].$$

Plugging the last estimate into (4.7.27) delivers

$$q^{k_l} \geq \mathbb{P}[\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}) > -\upsilon] \tag{4.7.28}$$

for large enough l. Letting $l \to \infty$, and with (4.7.25) in mind, we get $q \geq 1$.

Case 3. If $\alpha \in \mathbb{R}$ is such that $\alpha < \mathcal{I}_t(\lambda) - K_1$, we can proceed with the same discussion as in *Case 2.* The only difference is that the estimates analogous to (4.7.27) and (4.7.28) will involve the first queues in the sequence of tandems instead of the second ones. Again, the outcome is $q \geq 1$.

Case 4. The final case of $\alpha = \mathcal{I}_t(\lambda) - K_1$ is the most interesting one and will require consideration of three subcases. We start with some more notation. Let $\nu_l = \sqrt{n_{k_l}}(K_1 - \mathcal{I}_t(\lambda) + \alpha_{n_{k_l}})$,

$$Y^{(1,l)} = \sqrt{n_{k_l}}\,\bar{X}_t^{(1,n_{k_l})}(\mu_{n_{k_l}}) \quad \text{and} \quad Y^{(2,l)} = \sqrt{n_{k_l}}\,\bar{X}_t^{(2,n_{k_l})}(\mu_{n_{k_l}}).$$

According to the Central Limit Theorem,

$$Y^{(1,l)} \Rightarrow Y \quad \text{and} \quad Y^{(2,l)} \Rightarrow Z, \text{ as } l \to \infty, \tag{4.7.29}$$

where Y is a normal random variable with mean zero and variance

$$\sigma_Y^2 = \mathcal{I}_t(\lambda) + \lim_{l\to\infty} \alpha_{n_{k_l}} = \mathcal{I}_t(\lambda) + \alpha,$$

, and Z is also normal with mean zero, but with variance

$$\sigma_Z^2 = \lim_{l\to\infty} \alpha_{n_{k_l}} + \mathcal{I}_t(\mu^2) = \alpha + \mathcal{I}_t(\mu^2).$$

Since $\bar{X}_t^P = \mathcal{I}_t(\lambda - \mu^2) = K_1 + K_2$, we get that $\sigma_Z^2 = \mathcal{I}_t(\lambda) + \alpha - (K_1 + K_2)$.

Let $\{\nu_{l_i}\}$ be a subsequence of $\{\nu_l\}$ converging to a value $\nu \in \bar{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$, and let $\{r^i\}$ denote the subsequence $\{q^{k_{l_i}}\}$. It obviously converges to q. In this notation, we have

$$r^i = \mathbb{P}[Y^{(1,l_i)} > \nu_{l_i}] + \mathbb{P}[Y^{(2,l_i)} > -\nu_{l_i}], \text{ for all } i. \tag{4.7.30}$$

Below are the three subcases based on the value of ν.

a. $\nu = -\infty$
   Let M be an arbitrary positive constant. Then for all sufficiently large i, $\nu_{l_i} < -M$, and

$$r^i \geq \mathbb{P}[Y^{(1,l_i)} > -M].$$

   Employing (4.7.29), we conclude that $q = \lim_{i\to\infty} r^i \geq F_Y(M)$, with F_Y being the distribution function of Y. Arbitrariness of M allows us to seek the limit as $M \to \infty$ and get $q \geq 1$.

b. $\nu = +\infty$
   If we went ahead with the same reasoning as in the previous case, but applied to the processes derived from the second queues in the sequence of tandems, we would reach the same conclusion, i.e., $q \geq 1$.

c. $\nu \in \mathbb{R}$
   Since the sequence $\{\nu_{l_i}\}_i$ converges to the deterministic value ν, and the weak convergence in (4.7.29) holds true, we can let $i \to \infty$ in (4.7.30) and obtain

$$q = \mathbb{P}[Y > \nu] + \mathbb{P}[Z > -\nu].$$

Backtracking through all the stated cases, we get

$$q \geq 1 \wedge (\mathbb{P}[Y > \nu] + \mathbb{P}[Z > -\nu]).$$

In order to explore this lower bound further, we introduce $g : \mathbb{R} \to \mathbb{R}$ defined as $g(x) = \mathbb{P}[Y > x] + \mathbb{P}[Z > -x]$. Is there an absolute minimum of this function? The answer to this question is nothing but a simple exercise in calculus.

First, let us rewrite g as $g(x) = \mathbb{P}[Y < -x] + \mathbb{P}[Z < x]$. The first derivative of g can be written as

$$g'(x) = -f_Y(-x) + f_Z(x), \tag{4.7.31}$$

with f_Y and f_Z being the probability density functions of Y and Z, respectively. A suitable form for the derivative is obtained as

$$\begin{aligned} g'(x) &= -\frac{1}{\sqrt{2\pi}\sigma_Y} e^{-\frac{x^2}{2\sigma_Y^2}} + \frac{1}{\sqrt{2\pi}\sigma_Z} e^{-\frac{x^2}{2\sigma_Z^2}} \\ &= \frac{1}{\sqrt{2\pi}\sigma_Y} e^{-\frac{x^2}{2\sigma_Y^2}} \left(-1 + \frac{\sigma_Y}{\sigma_Z} \exp\left\{ -\frac{x^2}{2} \left(\frac{1}{\sigma_Z^2} - \frac{1}{\sigma_Y^2} \right) \right\} \right). \end{aligned}$$

Solving for x in $g'(x) = 0$ leads us to the equation

$$\frac{\sigma_Y}{\sigma_Z} \exp\left\{ -\frac{x^2}{2} \left(\frac{1}{\sigma_Z^2} - \frac{1}{\sigma_Y^2} \right) \right\} = 1$$

and its simpler equivalent form

$$x^2 = \frac{2\sigma_Y^2\sigma_Z^2}{K_1 + K_2} \ln\left(\frac{\sigma_Y}{\sigma_Z} \right). \tag{4.7.32}$$

Bearing in mind that $\sigma_Y > \sigma_Z$, we see that there exists a pair of solutions to this equation $\pm x_*$, with $x_* > 0$. Furthermore, for all x such that $x^2 > x_*^2$, $g'(x) < 0$ and for all x such that $x^2 < x_*^2$, $g'(x) > 0$. Therefore, a local minimum is attained at $-x_*$. Let us find bounds on this value.

$$\begin{aligned} g(-x_*) &= \mathbb{P}[Y < x_*] + \mathbb{P}[Z < -x_*] \\ &= 1 - \mathbb{P}[Y \geq x_*] + \mathbb{P}[Z > x_*] \\ &= 1 - \int_{x_*}^{\infty} (f_Y(x) - f_Z(x))\, dx. \end{aligned}$$

Because f_Y is a density of a normal random variable centered at zero and, hence, an even function, we have

$$g(-x_*) = 1 - \int_{x_*}^{\infty} (f_Y(-x) - f_Z(x))\, dx.$$

Recalling that the for all x such that $x^2 > x_*^2$, we have that $g'(x) = -(f_Y(-x) - f_Z(x)) < 0$, we conclude that

$$g(-x_*) < 1.$$

Taking the limit $\lim_{x\to\infty} g(x) = 1$, we realize that $g(-x_*)$ is the global minimum. On the other hand, since $x_* > 0$,

$$g(-x_*) = \mathbb{P}[Y < x_*] + \mathbb{P}[Z < -x_*] \geq \mathbb{P}[Y < x_*] > \frac{1}{2}.$$

This bound holds regardless of the choice of t; therefore, as long as $\mathrm{meas}(A) > 0$, we have

$$\int_A \liminf_{n\to\infty} p_t^n(\mu_n)\, dt > \frac{1}{2}\mathrm{meas}(A).$$

Recalling the sufficient condition (4.7.26), we see that if, in addition to Assumption 4.7.1, the parameters λ and μ^2 satisfy the following assumption, it is impossible to approach the asymptotically optimal performance of sequences of stochastic disciplines using a sequence of deterministic disciplines.

Assumption 4.7.3. Let λ and μ^2 satisfy

$$\mathrm{meas}\{t \in [0,T] : \bar{X}_t^P = K_1 + K_2\} > 0.$$

Finally, let us provide an example of a set of parameters for the tandem system and a sequence of deterministic service disciplines such that the above lower bound is achieved.

Example 4.7.4. Let $T = 1$, $m = \infty$ and $K_1 = K_2 = \frac{1}{3}$. Let the arrival rate in the first station be $\lambda = \mathbf{1}_{[0,\frac{2}{3}]}$ and let there be no service in the second station at all, i.e., let $\mu^2 \equiv 0$.

This set of parameters obviously satisfies Assumption 4.7.1. The length of the fluid pooled queue associated with this system equals $\bar{Q}_t^P = \bar{X}_t^P = \min(t, \frac{2}{3})$. Hence, Assumption 4.7.3 is also satisfied.

With (4.7.32) in mind, we define the function $\nu : [\frac{2}{3}, T] \to \mathbb{R}$ by

$$\nu_t = -\sqrt{\ln\left(\frac{2\mathcal{I}_t(\lambda) - K_1}{2\mathcal{I}_t(\lambda) - 2K_1 - K_2}\right)\frac{(2\mathcal{I}_t(\lambda) - K_1)(2\mathcal{I}_t(\lambda) - 2K_1 - K_2)}{K_1 + K_2}} = -\sqrt{\frac{\ln(3)}{2}}.$$

We see that ν is a constant function. So let us introduce $\nu^* = \sqrt{\frac{\ln(3)}{2}}$.

Next, we define a sequence of deterministic service discipline in the following manner. For all n, let $\nu^n = \frac{\nu^*}{\sqrt{n}}$ and $\mu_n = \frac{1}{2}\mathbf{1}_{[\nu^n,\frac{2}{3}]}$. As $n \to \infty$, we have $\mathcal{I}(\mu_n) \to \mathcal{I}(\mu^*)$, uniformly, with $\mu^* = \frac{1}{2}\mathbf{1}_{[0,\frac{2}{3}]}$. By Theorem 1.2.5 coupled with the lemma on p.151 from [Bil99], we infer that

$$\frac{1}{n}X^{(1,n)}(\mu_n) \to \bar{X}^{(1)}(\mu^*),$$

with $\bar{X}_t^{(1)}(\mu^*) = \frac{t}{2} \wedge \frac{1}{3}$. The continuity of the Skorokhod map implies

$$\frac{1}{n}Q^{(1,n)}(\mu_n) \to \Gamma(\bar{X}^{(1)}(\mu^*)) = \bar{X}^{(1)}(\mu^*). \tag{4.7.33}$$

As for the sequence of second queues in the sequence of tandem systems, we get by the same token

$$\frac{1}{n}X^{(2,n)}(\mu_n) \to \bar{X}^{(2)}(\mu^*),$$

for $\bar{X}^{(2)}_t(\mu^*) = \frac{t}{2} \wedge \frac{1}{3}$, and

$$\frac{1}{n}Q^{(2,n)}(\mu_n) \to \Gamma(\bar{X}^{(2)}(\mu^*)) = \bar{X}^{(2)}(\mu^*). \tag{4.7.34}$$

Inheriting the definitions of sets A, B and C from (4.7.10), and seeing that $\bar{X}^P_t = t \wedge \frac{2}{3}$, we conclude that $A = [\frac{2}{3}, 1]$, $B = [0, \frac{2}{3})$ and $C = \emptyset$. The expected penalty in the n^{th} system becomes

$$\begin{aligned}
&\int_0^T \left(\mathbb{P}\left[\frac{1}{n}Q^{(1,n)}_t(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}Q^{(2,n)}_t(\mu_n) > K_2\right]\right) dt \\
&\quad = \int_A \left(\mathbb{P}\left[\frac{1}{n}Q^{(1,n)}_t(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}Q^{(2,n)}_t(\mu_n) > K_2\right]\right) dt \\
&\qquad\quad + \int_B \left(\mathbb{P}\left[\frac{1}{n}Q^{(1,n)}_t(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}Q^{(2,n)}_t(\mu_n) > K_2\right]\right) dt.
\end{aligned}$$

Let us focus on the region B first. Thanks to (respectively) (4.7.33) and (4.7.34), for any $t \in B$, we have

$$\frac{1}{n}Q^{(1,n)}_t(\mu_n) \to \bar{X}^{(1)}_t(\mu^*) < \frac{1}{3} = K_1$$

and

$$\frac{1}{n}Q^{(2,n)}_t(\mu_n) \to \bar{X}^{(2)}_t(\mu^*) < \frac{1}{3} = K_2.$$

Thus, by Lebesgue's Dominated Convergence Theorem

$$\lim_{n\to\infty} \int_B \left(\mathbb{P}\left[\frac{1}{n}Q^{(1,n)}_t(\mu_n) > K_1\right] + \mathbb{P}\left[\frac{1}{n}Q^{(2,n)}_t(\mu_n) > K_2\right]\right) dt = 0.$$

Evaluating the limiting penalty aggregated over region A requires second order convergence results for the sequence of tandems. For all $t \in A$, we first bound the first probability in the integrand above in a natural way as

$$\begin{aligned}
\mathbb{P}\left[\frac{1}{n}Q^{(1,n)}_t(\mu_n) > K_1\right] &\geq \mathbb{P}\left[\frac{1}{n}X^{(1,n)}_t(\mu_n) > K_1\right] \\
&= \mathbb{P}\left[\sqrt{n}(\frac{1}{n}X^{(1,n)}_t(\mu_n) - \mathcal{I}_t(\lambda - \mu_n)) - \sqrt{n}(\mathcal{I}_t(\lambda - \mu_n) - K_1) > 0\right].
\end{aligned}$$

Using the definition of the sequence $\{\mu_n\}$, we further obtain that

$$\mathbb{P}\left[\frac{1}{n}X_t^{(1,n)}(\mu_n) > K_1\right] = \mathbb{P}\left[\sqrt{n}(\frac{1}{n}X_t^{(1,n)}(\mu_n) - \mathcal{I}_t(\lambda - \mu_n)) > \nu^*\right].$$

However, by the Central Limit Theorem, the last probability converges as follows

$$\mathbb{P}\left[\sqrt{n}(\frac{1}{n}X_t^{(1,n)}(\mu_n) - \mathcal{I}_t(\lambda - \mu_n)) > \nu^*\right] \to 1 - F_N\left(\frac{\nu^*}{\sqrt{\mathcal{I}_t(\lambda + \mu^*)}}\right),$$

where F_N denotes the standard normal distribution function. We can further evaluate this limit using the explicit definitions of the rates involved and obtain

$$1 - F_N\left(\frac{\nu^*}{\mathcal{I}_t(\lambda + \mu^*)}\right) = 1 - F_N(\nu^*) = 1 - F_N\left(\sqrt{\frac{\ln(3)}{2}}\right) = F_N\left(-\sqrt{\frac{\ln(3)}{2}}\right).$$

By the same argument, we get that

$$\mathbb{P}\left[\frac{1}{n}Q_t^{(2,n)}(\mu_n) > K_2\right] \geq \mathbb{P}\left[\frac{1}{n}X_t^{(2,n)}(\mu_n) > K_2\right] \to F_N(\nu^*) = F_N\left(\sqrt{\frac{\ln(3)}{2}}\right).$$

Both limits are independent of the instant t. Therefore, on the segment A, which is of length $\frac{1}{3}$, the aggregated penalty equals $\frac{1}{3}(F_N(\sqrt{\frac{\ln(3)}{2}}) + F_N(-\sqrt{\frac{\ln(3)}{2}})) \approx 0.26$.

Let us now concentrate on the performance of the corresponding pooled queue. This queue is such that its fluid-limit is strictly below $K_1 + K_2$ for all $t < \frac{2}{3}$. On the other hand, for every $t \geq \frac{2}{3}$, the centered and scaled queue length, by the Central Limit Theorem converges in distribution to a normal random variable centered at zero. Over the period $[\frac{2}{3}, 1]$, we have that

$$\mathbb{P}[\frac{1}{n}X_t^{(P,n)} > K_1 + K_2] \to F_N(0) = \frac{1}{2}.$$

The limit of the expected amount of time the pooled queue spends above the level $K_1 + K_2$, hence, equals $\frac{1}{6}$. This value is strictly less than the minimal limiting penalty possible for the tandem system as evaluated above.

Remark 4.7.1. Taking a closer look at the construction of a particular fluid-optimal discipline in Subsection 4.4.4, we see that the discipline μ^* of (4.4.11) and μ^* of the last example coincide for the particular choice of parameters in the last example.

4.8 Asymptotic Optimality - Finite Buffers

4.8.1 Reduction to a Sufficient Subclass of Admissible Sequences

We first commit our attention to a simplification of the class of admissible disciplines it suffices to consider in our optimization problem. The flow of reasoning in the sequel is analogous to the one in Lemma 4.5.1 and its proof. While all the statements are plausible, the model and the notion of control in the stochastic case are more complex and their rigorous treatment requires a far more detailed approach.

In what follows, we will suppress the index of uniform acceleration n for the sake of simplicity. All the statements remain valid regardless of the choice of n, and their proofs carry over word-for-word. In the same spirit, for all $\mu \in \mathcal{B}$, where $\mathcal{B}$ is defined in (B.3.5), we neglect the index $n = 1$ and look at the following queue lengths

$$\begin{aligned} Q^{(1)}(\mu) &= N_1^+(\mathcal{I}(\lambda)) - N_1^-(\mathcal{I}(\mu)) + L^{(1)}(\mu) - U^{(1)}(\mu), \\ Q^{(2)}(\mu) &= N_1^-(\mathcal{I}(\mu)) - L^{(1)}(\mu) - N_2^-(\mathcal{I}(\mu^2)) + L^{(2)}(\mu) - U^{(2)}(\mu). \end{aligned} \tag{4.8.1}$$

All the processes in the above display are interpreted in the fashion consistent with the representation of the two-sided reflection maps Γ^{K_1} and Γ^{K_2} of Definition 1.2.2 applied to $N_1^+(\mathcal{I}(\lambda)) - N_1^-(\mathcal{I}(\mu))$ and $N_1^-(\mathcal{I}(\mu)) - L^{(1)}(\mu) - N_2^-(\mathcal{I}(\mu^2))$, respectively.

Lemma 4.8.1. *Consider an admissible service discipline $\mu \in \mathcal{B}$. Then there exists a service discipline $\tilde{\mu} \in \mathcal{B}$ such that the following relations hold almost surely:*

1. $L^{(1)}(\tilde{\mu}) \equiv 0$;
2. $U^{(2)}(\tilde{\mu}) \equiv 0$;
3. $U^{(1)}(\tilde{\mu}) \leq U^{(1)}(\mu) + U^{(2)}(\mu)$.

Proof. See Appendix D.5 □

Based on the above lemma, from now on we consider only admissible service disciplines from the following space

$$\mathcal{B}^* = \{\mu \in \mathcal{B} :\ L^{(1)}(\mu) \equiv 0 \,\text{and}\, U^{(2)}(\mu) \equiv 0\}. \tag{4.8.2}$$

4.8.2 Lower Bound

We proceed with the definition of the pooled queue with the finite capacity $K = K_1 + K_2$, corresponding to the tandem system described in (4.8.1). The netput process generating the queue length in this setting equals

$$X^P = N_1^+(\mathcal{I}(\lambda)) - N_2^-(\mathcal{I}(\mu^2)). \tag{4.8.3}$$

Once the two-sided regulator is applied to the process X^P, we obtain the length of the pooled queue and write it in the familiar way using the minimal lower and upper regulator processes, i.e., we have

$$Q^P = N_1^+(\mathcal{I}(\lambda)) - N_2^-(\mathcal{I}(\mu^2)) + L^P - U^P. \tag{4.8.4}$$

The next lemma contains the contribution of the pooled queue to a convenient lower bound. Its proof is a *mutatis mutandis* repetition of the proof of Lemma 4.5.2, so we exhibit it in Appendix D.5.

Lemma 4.8.2. *For all $\mu \in \mathcal{B}^*$, we have that*

(i) $L^P \leq L^{(2)}(\mu)$, *almost surely.*

(ii) $U^{(1)}(\mu) \geq U^P$, *almost surely.*

The following is a useful consequence of Lemmas 4.8.1 and 4.8.2.

Corollary 4.8.3. *For all $\mu \in \mathcal{B}$, we have that $U^{(1)}(\mu) + U^{(2)}(\mu) \geq U^P$, almost surely.*

We have learned in the study of the fluid limit that the comparison with the pooled queue is informative only in the region before the constraint on the total amount of service causes necessary upper regulation in the first queue. Hence, we present the following lemma.

Lemma 4.8.4. *For every $\mu \in \mathcal{B}$ and every $T' \leq T$, we have that*

$$U_T^{(1)}(\mu) + U_T^{(2)}(\mu) \geq U_{T'}^P \vee [N_1^+(\mathcal{I}_T(\lambda)) - N_1^-(m) - K_1]^+, \textit{ a.s.}$$

Proof. As usual, let us fix μ and T'. According to Lemma 4.8.1, there exists a $\tilde{\mu} \in \mathcal{B}^*$, such that

$$U^{(1)}(\mu) + U^{(2)}(\mu) \geq U^{(1)}(\tilde{\mu}), \textit{ a.s.}$$

Then we have the following chain of inequalities, all holding in the almost sure sense:

$$U_T^{(1)}(\tilde{\mu}) \geq U_{T'}^{(1)}(\tilde{\mu}) \vee \sup_{s \in (T',T]} [N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\tilde{\mu})) - K_1]^+.$$

By Lemma 4.8.2, the last inequality yields

$$U_T^{(1)}(\tilde{\mu}) \geq U_{T'}^P \vee \sup_{s \in (T',T]} [N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\tilde{\mu})) - K_1]^+.$$

Finally, using the constraint on total service available and the increase of Poisson processes, we get

$$U_T^{(1)}(\tilde{\mu}) \geq U_{T'}^P(\tilde{\mu}) \vee [N_1^+(\mathcal{I}_T(\lambda)) - N_1^-(m) - K_1]^+.$$

□

4.8.3 An Asymptotic Lower Bound

Next, we need to explore more thoroughly the limiting behavior of the sequence of pooled queues $\{Q^{(P,n)}\}$, where $Q^{(P,n)} = \Gamma^{nK}(X^{(P,n)})$, for the mapping Γ^{nK} from Definition 1.2.2, and

$$X^{(P,n)} = N_1^+(\mathcal{I}(n\lambda)) - N_2^-(n\mathcal{I}(\mu^2)).$$

Theorem 1.2.5 applied to the sequence of netput processes $\{X^{(P,n)}\}$ in the pooled system, yields

$$\frac{1}{n}X^{(P,n)} \to \bar{X}^P = \mathcal{I}(\lambda - \mu^2),\ a.s.,$$

uniformly. By Proposition 1.2.4, we conclude that

$$\begin{aligned} &\frac{1}{n}Q^{(P,n)} \to \bar{Q}^P, \\ &\frac{1}{n}L^{(P,n)} \to \bar{L}^P, \\ &\frac{1}{n}U^{(P,n)} \to \bar{U}^P, \end{aligned} \tag{4.8.5}$$

almost surely, uniformly on compacts, where $\bar{Q}^P = \Gamma^K(\bar{X}^P)$ and $\bar{L}^P$ and $\bar{U}^P$ are associated with $\bar{X}^P$ and K in the sense of Definition 1.2.2. All the statements in the previous subsection regarding the subclass of controls sufficient for optimality and the lower bound on the performance is easily generalized to the uniformly accelerated sequence of systems. For every n, we refer to the sufficient class of controls as

$$\mathcal{B}_n^* = \{\mu \in \mathcal{L}^{(n)}(m) : L^{(1,n)}(\mu) \equiv 0 \text{ and } U^{(2,n)}(\mu) \equiv 0\}. \tag{4.8.6}$$

Lemma 4.8.5. *For every admissible sequence $\{\mu_n\}$, we have that*

$$\liminf_{n\to\infty} \mathbb{E}[J_F^{(n)}(\mu_n) - J_{FLB}^{(n)}] \geq 0,$$

where

$$J_{FLB}^{(n)} = \frac{1}{n}U_{T^*}^{(P,n)} \vee \left[\frac{1}{n}N_1^+(n\mathcal{I}_T(\lambda)) - \frac{1}{n}N_1^-(nm) - K_1\right]^+ \tag{4.8.7}$$

with

$$T^* = \tau^*(\mu^{*,F}) = \inf\{t > 0 : \mathcal{I}_t(\mu^{*,F}) = m\} \wedge T, \tag{4.8.8}$$

and where $\mu^{*,F}$ is the fluid-optimal deterministic service discipline given in (4.5.37).

Proof. From Lemma 4.8.1, we know that it suffices to consider only admissible sequences $\{\mu_n\}$ such that $\mu_n \in \mathcal{B}_n^*$, where the subclass $\mathcal{B}_n^*$ is defined in (4.8.6). Let us fix such an admissible sequence $\{\mu_n\}$. For every n, the performance of the service discipline μ_n can be rewritten and bounded from below in the following fashion:

$$J_F^{(n)}(\mu_n) = \frac{1}{n}U_T^{(1,n)}(\mu_n) + \frac{1}{n}U_T^{(2,n)}(\mu_n) \geq \frac{1}{n}U_T^{(1,n)}(\mu_n).$$

By the second equality in (1.2.6), we can rewrite the lower bound in the last display as

$$\begin{aligned}
J_F^{(n)}(\mu_n) &\geq \frac{1}{n}\sup_{s\in[0,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) + L_s^{(1,n)}(\mu_n) - nK_1]^+ \\
&\geq \frac{1}{n}\sup_{s\in[0,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) - nK_1]^+ \\
&= \frac{1}{n}\Big(\sup_{s\in[0,T^*]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) - nK_1]^+ \\
&\qquad \vee \sup_{s\in(T^*,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) - nK_1]^+\Big) \\
&= \frac{1}{n}\left(U_{T^*}^{(1,n)}(\mu_n) \vee \sup_{s\in(T^*,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) - nK_1]^+\right).
\end{aligned}$$

By Corollary 4.8.3 and the last display, we have that

$$J_F^{(n)}(\mu_n) \geq \frac{1}{n}\left(U_{T^*}^{(P,n)} \vee \sup_{s\in(T^*,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu_n)) - nK_1]^+\right).$$

Due to the constraint $\mathcal{I}_T(\mu_n) \leq m$ and the monotonicity of the Poisson process N_1^-, we deduce that

$$J_F^{(n)}(\mu_n) \geq \frac{1}{n}\left(U_{T^*}^{(P,n)} \vee \sup_{s\in(T^*,T]}[N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(nm) - nK_1]^+\right).$$

The above sequence of inequalities evidently implies the announced result. □

4.8.4 A Particular Admissible Sequence

Recall the deterministic service discipline $\mu^{*,F}$ defined in (4.5.24) and set $\mu_n = \mu^{*,F}$, for every n. Since all the terms in the sequence $\{\mu_n\}$ are deterministic, they are trivially appropriately adapted to the filtration constructed in Appendix B.3. Moreover, the constraint on the total amount of service available is straightforwardly satisfied by the definition of $\mu^{*,F}$.

For every n, the lengths of the queues in tandem with $\mu_n = \mu^{*,F}$ as the service discipline are given by

$$\begin{aligned} Q^{(1,n)}(\mu^{*,F}) &= N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu^{*,F})) + L^{(1,n)}(\mu^{*,F}) - U^{(1,n)}(\mu^{*,F}), \\ Q^{(2,n)}(\mu^{*,F}) &= N_1^-(n\mathcal{I}(\mu^{*,F})) - L^{(1,n)}(\mu^{*,F}) - N_2^-(n\mathcal{I}(\mu^2)) \\ &\qquad + L^{(2,n)}(\mu^{*,F}) - U^{(2,n)}(\mu^{*,F}). \end{aligned} \tag{4.8.9}$$

The netput process in the first station equals

$$X^{(1,n)}(\mu^{*,F}) = N_1^+(n\mathcal{I}(\lambda)) - N_1^-(n\mathcal{I}(\mu^{*,F})) \tag{4.8.10}$$

for every n. By Proposition 1.2.4, we conclude that as $n \to \infty$

$$\begin{aligned} \frac{1}{n} Q^{(1,n)}(\mu^{*,F}) &\to \bar{Q}^{(1)}(\mu^{*,F}), \\ \frac{1}{n} L^{(1,n)}(\mu^{*,F}) &\to \bar{L}^{(1)}(\mu^{*,F}), \\ \frac{1}{n} U^{(1,n)}(\mu^{*,F}) &\to \bar{U}^{(1)}(\mu^{*,F}), \end{aligned} \tag{4.8.11}$$

almost surely, in the uniform topology, where $\bar{Q}^{(1)}(\mu^{*,F}) = \Gamma^K(\bar{X}^{(1)}(\mu^{*,F}))$, and $\bar{L}^{(1)}(\mu^{*,F})$ and $\bar{U}^{(1)}(\mu^{*,F})$ are the lower and upper regulators associated with $\bar{X}^{(1)}(\mu^{*,F})$ and K_1 in the sense of Definition 1.2.2. Combining (4.8.11) with (4.5.31), we conclude that as $n \to \infty$,

$$\frac{1}{n} L^{(1,n)}(\mu^{*,F}) \to 0, \ a.s., \tag{4.8.12}$$

uniformly. The last display, in conjunction with Theorem 1.2.5, yields that as $n \to \infty$

$$\frac{1}{n}\left(N_1^-(n\mathcal{I}(\mu^{*,F})) - N_2^-(n\mathcal{I}(\mu^2)) - \bar{L}^{(1)}(\mu^{*,F})\right) \to \mathcal{I}(\mu^{*,F} - \mu^2), \ a.s., \tag{4.8.13}$$

uniformly. Thanks to Proposition 1.2.4 and (4.5.35), we have

$$\begin{aligned} \frac{1}{n} Q^{(2,n)}(\mu^{*,F}) &\to \bar{Q}^{(2)}(\mu^{*,F}), \\ \frac{1}{n} L^{(2,n)}(\mu^{*,F}) &\to \bar{L}^{(2)}(\mu^{*,F}), \\ \frac{1}{n} U^{(2,n)}(\mu^{*,F}) &\to \bar{U}^{(2)}(\mu^{*,F}) \equiv 0, \end{aligned} \tag{4.8.14}$$

almost surely, in the uniform topology.

Lemma 4.8.6. *As $n \to \infty$,*

$$\frac{1}{n}\left(U^{(1,n)}(\mu^{*,F}) - U^{(P,n)}\right) \to 0, \ a.s., \tag{4.8.15}$$

in the uniform topology on the segment $[0, T^*]$, where T^* is given in (4.8.8).

Proof. From the third limit in display (4.8.11) and Lemma 4.5.3, we conclude that

$$\frac{1}{n} U^{(1,n)}(\mu^{*,F}) \to \bar{U}^P, \ a.s., \tag{4.8.16}$$

uniformly on $[0, T^*]$. Combining (4.8.16) with the third claim in (4.8.5), which is made possible by the continuity of addition in the uniform topology, we obtain the desired result. □

A direct consequence of the last lemma is the limiting value of the performance of the admissible sequence whose terms are identically equal to $\mu^{*,F}$, up to the deterministic time T^*.

Corollary 4.8.7. *As $n \to \infty$, we have that*

$$\frac{1}{n}\left(U^{(1,n)}_{T^*}(\mu^{*,F}) + U^{(2,n)}_{T^*}(\mu^{*,F}) - U^{(P,n)}_{T^*}\right) \to 0, \ a.s.$$

Proof. Simple addition of the third limit in (4.8.14) at time T^*, and the result of Lemma 4.8.6 yields the announced result. □

Next, we evaluate the limiting penalty incurred by the service $\mu^{*,F}$ after time T^*.

Lemma 4.8.8. *As $n \to \infty$, the following holds true in the almost sure sense:*

$$\frac{1}{n}\left[U^{(1,n)}_{T}(\mu^{*,F}) - U^{(P,n)}_{T^*} \vee [N_1^+(n\mathcal{I}_T(\lambda)) - N_1^-(nm) - nK_1]^+\right] \to 0. \tag{4.8.17}$$

Proof. For every n, the total amount of downward pushing in the first station when $\mu^{*,F}$ is used as the service discipline is

$$U^{(1,n)}_{T}(\mu^{*,F}) = U^{(1,n)}_{T^*}(\mu^{*,F}) \vee \sup_{s\in(T^*,T]} [N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(n\mathcal{I}_s(\mu^{*,F})) - nK_1]^+.$$

By the definition of T^*, the above equality can be transformed to

$$U^{(1,n)}_{T}(\mu^{*,F}) = U^{(1,n)}_{T^*}(\mu^{*,F}) \vee \sup_{s\in(T^*,T]} [N_1^+(n\mathcal{I}_s(\lambda)) - N_1^-(nm) - nK_1]^+.$$

Due to the increase of the Poisson process N_1^+, we further get that

$$U^{(1,n)}_{T}(\mu^{*,F}) = U^{(1,n)}_{T^*}(\mu^{*,F}) \vee [N_1^+(n\mathcal{I}_T(\lambda)) - N_1^-(nm) - nK_1]^+.$$

By Lemma D.5.1 and the last display, the left-hand side of the expression (4.8.17) can be bounded from above by

$$\frac{1}{n}|U^{(1,n)}_{T^*}(\mu^{*,F}) - U^{(P,n)}_{T^*}|.$$

However, this term has limit zero thanks to Lemma 4.8.6. □

Proposition 4.8.9. *As $n \to \infty$,*

$$J_F^{(n)}(\mu^{*,F}) - J_{FLB}^{(n)} \to 0, \ a.s.,$$

where $J_{FLB}^{(n)}$ is given by (4.8.7).

Proof. Thanks to the third claim in (4.8.14), we can discard the limiting contribution of the upper regulator in the second queue in the tandem. Hence, as $n \to \infty$,

$$J_F^{(n)}(\mu^{*,F}) - \frac{1}{n} U_T^{(1,n)}(\mu^{*,F}) \to 0, \ a.s.$$

Combining this limit with Lemma 4.8.8 and the third claim in (4.8.5), we obtain the desired convergence. □

4.8.5 Asymptotic Optimality

Finally, we complete the task of formulating and solving the asymptotic control problem in the finite-buffer setting for the tandem system.

Definition 4.8.10. An admissible sequence $\{\mu_n^*\}$ is called *asymptotically optimal* for the sequence of performance measures $\{J_F^{(n)}\}$, if

$$\liminf_{n\to\infty} \mathbb{E}[J_F^{(n)}(\mu_n) - J_F^{(n)}(\mu_n^*)] \geq 0,$$

for any other admissible sequence $\{\mu_n\}$.

Theorem 4.8.11. *The sequence identically equal to $\mu^{*,F}$ is asymptotically optimal for performance measures $\{J_F^{(n)}\}$.*

Proof. Let $\{\mu_n\}$ be any admissible sequence, then we have that

$$\begin{aligned}&\liminf_{n\to\infty} \mathbb{E}[J_F^{(n)}(\mu_n) - J_F^{(n)}(\mu^{*,F})] \\ &\qquad \geq \liminf_{n\to\infty} \mathbb{E}[J_F^{(n)}(\mu_n) - J_{FLB}^{(n)}] + \lim_{n\to\infty} \mathbb{E}[J_{FLB}^{(n)} - J_F^{(n)}(\mu^{*,F})].\end{aligned}$$

Using Lemma 4.8.5 and Proposition 4.8.9, along with boundedness of all random variables involved, we complete the proof. □

Appendix A

A General Optimization Problem

Here we formulate an abstract control problem and discuss the manner in which it generates the particular control problems analyzed throughout this document.

A.1 A General Control Problem

Let us return to the general single-station model described in the first paragraph of Subsection 3.1.1. We do not even consider the system with finite capacity. Our only assumptions are that the arrival process is Poisson with a known rate and that the potential service process is also Poisson with a rate subject to our control. The natural ambient space for the service rate is the space of all nonnegative integrable functions on $[0, T]$, denoted as before by $\mathbb{L}^1_+[0, T]$.

Next, we introduce two mappings that encapsulate the cost structure. The mapping $J : \mathbb{L}^1_+[0, T] \to \mathbb{R}_+$ denotes the penalty incurred by a service discipline. The convex, nondecreasing function $c : \mathbb{R}_+ \to \mathbb{R}_+$ represents the cost associated with the instantaneous service rate, i.e., the total cost of a chosen service $\mu \in \mathbb{L}^1_+[0, T]$ is given by

$$C(\mu) = \int_0^T c(\mu_s)\, ds. \tag{A.1.1}$$

The expression in (A.1.1) implicitly defines the mapping $C : \mathbb{L}^1_+[0, T] \to \mathbb{R}_+$ as the aggregate cost of service. Hence, for every $\mu \in \mathbb{L}^1_+[0, T]$ its measure of performance is the value of $J(\mu) + C(\mu)$. We wish to minimize this value on the set $\mathbb{L}^1_+[0, T]$, i.e., we seek the value J^*, where

$$J^* = \inf_{\mu \in \mathbb{L}^1_+[0,T]} \{J(\mu) + C(\mu)\}. \tag{A.1.2}$$

Lemma A.1.1. *With mappings J and C as above, the following equality holds true*

$$J^* = \inf_{M \geq 0} \inf\{J(\mu) + M : \mu \in \mathbb{L}^1_+[0, T],\, C(\mu) \leq M\}. \tag{A.1.3}$$

Proof. First, we rewrite J^* as

$$\begin{aligned} J^* &= \inf_{M\geq 0} \inf\{J(\mu) + C(\mu) : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) = M\} \\ &= \inf_{M\geq 0} \inf\{J(\mu) + M : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) = M\} \\ &= \inf_{M\geq 0} \left(\inf\{J(\mu) : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) = M\} + M\right). \end{aligned} \tag{A.1.4}$$

This calculation yields that, in order to prove the lemma, it is enough to prove that

$$\begin{aligned} &\inf_{M\geq 0} \inf\{J(\mu) + C(\mu) : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) = M\} \\ &\qquad = \inf_{M\geq 0} \inf\{J(\mu) + M : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) \leq M\}. \end{aligned} \tag{A.1.5}$$

Obviously, the left-hand side of (A.1.5) exceeds or is equal to the right-hand side of the same expression.

Let us suppose that the left-hand side is strictly greater than the right-hand side in (A.1.5). Then there necessarily exists a positive constant M' such that

$$J^* > \inf\{J(\mu) + M' : \mu \in \mathbb{L}^1_+[0,T],\, C(\mu) \leq M'\}.$$

Moreover, there also exists a nonnegative integrable function μ' satisfying $C(\mu) \leq M'$ and

$$J^* > J(\mu') + M'.$$

The last inequality contradicts the expression for J^* which we obtained in (A.1.4). □

A.2 Correspondence with the Current Control Problem

In the context of different versions of the control problem we are addressing, the general penalty function J corresponds either to the mappings counting the number of lost jobs in the finite buffer model, or to the mappings measuring the amount of time spent above the threshold in the infinite buffer model.

The cost of instantaneous service c is for us always the identity, which generates the total cost $\int_0^T \mu_s\, ds$ for a given service discipline μ.

Finally, we do not directly attempt to minimize (A.1.2). Instead, as a stepping stone to this problem, we restrict our attention to the study of the "inner infimum" in the expression (A.1.3). The given constraint on the cost (amount) of service is denoted in the main text by m.

Appendix B

On the Information Structure

B.1 The Filtration

We wish to allow for non-deterministic service disciplines, continuously adapting to the current state of the system, yet at the same time forbid any anticipation of the state of the system in the future.

Let us start with the part of flow of information not depending on the service discipline at all, i.e., which depends on the arrivals into the system only. For the Poisson process with rate of arrivals ρ^a, all the arrivals are reflected in the filtration generated by the arrival process $N^+(\mathcal{I}(\rho^a))$, where N^+ is a unit Poisson process. We start the formal construction with the filtration generated by N^+ in the usual manner as

$$\mathcal{F}_t^+ = \sigma\{s \leq t : N^+(s)\}, \text{ for all } t. \tag{B.1.1}$$

Of course, what we really need is the filtration generated by the time-changed Poisson process, so let us introduce filtration

$$\mathcal{G}_t^+ = \mathcal{F}^+_{\mathcal{I}_t(\rho^a)}, \text{ for all } t. \tag{B.1.2}$$

The filtration $\mathbb{G} = \{\mathcal{G}_t^+\}_{t\in[0,T]}$ gathers all the information on the arrivals in the system available to the controller.

The other standing assumption on our model is that the process of potential departures from the system behaves as a time-changed Poisson process as well. So, let as consider the Poisson point process ξ^d on the space $[0,T] \times [0,\infty)$ with the intensity measure given as the Lebesgue measure on the said space and denoted by ρ^d. Then, we can define the filtration containing the information gathered from the process of potential departures from the system as

$$\mathcal{F}_t = \sigma(\xi(A); A \in \mathcal{B}([0,t] \times [0,\infty))), \tag{B.1.3}$$

where $\mathcal{B}(Y)$ denotes the Borel $\sigma-$algebra of a metric space Y. The interpretation of this construction is that at each time t, we are able to establish the number of potential departures from the system as

$$\xi\{(s,x) : s \le t,\, x \le \mu_s\}.$$

We define the filtration $\{\mathcal{H}_t\}_{t\in[0,T]}$ as

$$\mathcal{H}_t = \mathcal{G}_t^+ \vee \mathcal{F}_t, \text{ for all } t \in [0,T]. \tag{B.1.4}$$

The filtration $\{\mathcal{H}_t\}_{t\in[0,T]}$ contains the information on the past and present events in the system of interest without anticipating the future of the system.

B.2 Flow of Information in the Single Station Problem

We now apply the findings of Appendix B.1 to the specific problem exhibited in Subsection 3.1.3. The following ingredients will be used to define the space of all admissible sequences of controls in the present problem:

(i) a sequence of independent unit Poisson processes $\{N^{(+,n)}\}$;

(ii) a sequence of independent Poisson point processes $\{\xi^n\}$ on $[0,T]\times[0,\infty)$ with the Lebesgue measure as the intensity measure and also independent of $\{N^{(+,n)}\}$.

For every index n, we modify the definition in (B.1.1) to get

$$\mathcal{F}_t^{(+,n)} = \sigma\{s \le t : N^{(+,n)}(s)\}, \text{ for all } t, \tag{B.2.1}$$

and in the manner of (B.1.2) obtain the desired filtration corresponding to the time-changed Poisson process

$$\mathcal{G}_t^{(+,n)} = \mathcal{F}_{n\mathcal{I}_t(\lambda)}^{(+,n)}, \text{ for all } t. \tag{B.2.2}$$

We set $\mathbb{G}^n = \{\mathcal{G}_t^{(+,n)}\}_{t\in[0,T]}$. The filtration $\mathbb{G}^n$ contains all the information on the arrivals into the n^{th} system in the sequence announced in Subsection 3.1.1.

Following (B.1.3), we define

$$\mathcal{F}_t^{(n)} = \sigma(\xi^n(A); A \in \mathcal{B}([0,t]\times[0,\infty))).$$

Finally, the filtration containing all the information available to the controller in the n^{th} system is given by

$$\mathcal{H}_t^{(n)} = \mathcal{G}_t^{(+,n)} \vee \mathcal{F}_t^{(n)} \qquad \text{for all } t \in [0,T]. \tag{B.2.3}$$

We say that a nonnegative random function μ on $[0,T]$ is an admissible service discipline in the n^{th} system if it is $\{\mathcal{H}_t^{(n)}\}-$predictable.

x

B.3 Information Flow in a More General Environment

Herein we set up the appropriate environment for the set of admissible disciplines referred to repeatedly in the main text. The challenge in doing so lies in reconciling the different time-scales at which different processes in a certain network develop. Moreover, as is evident from the models exhibited in Sections 4.1 and 4.3, the control-process itself governs the speed at which the information is acquired.

Although in this appendix we allow for more general network topologies than the ones discussed in the main text, we still assume that all the arrival and potential departure processes are modeled as time-inhomogeneous Poisson processes.

First, let us differentiate between the controlled and the uncontrolled processes in a network. Let the uncontrolled stochastic processes in a given finite network be denoted by $\{\mathcal{E}^j\}_{j=1}^J$. The information accumulated from these, exogenously governed random processes is gathered over the interval $[0,T]$ in the filtration $\mathbb{G} = \{\mathcal{G}_t\}_{t\in[0,T]}$ given by

$$\mathcal{G}_t = \sigma\{\mathcal{E}^j_s : s \leq t, 1 \leq j \leq J\}. \tag{B.3.1}$$

Let us denote the number of controlled processes in our finite network by V and and let us momentarily concentrate on a single controlled process indexed by some $v \in \{1, 2, \ldots, V\}$. The standing assumption is that this process is a time-changed Poisson process. So, let us introduce a family of independent Poisson point processes $\{\xi_1, \xi_2, \ldots, \xi_V\}$ on the space $[0,T] \times [0,\infty)$ and with the Lebesgue measure as the intensity measure. Regardless of the choice of control, the information gathered from the controlled random process is contained in the filtration

$$\mathcal{F}_t = \sigma(\xi_i(A); 1 \leq i \leq V,\ A \in \mathcal{B}([0,t] \times [0,\infty))). \tag{B.3.2}$$

Altogether, combining (B.3.1)and (B.3.2), all the information on the past and the present of the system is gathered in the filtration $\{\mathcal{H}_t\}$, where

$$\mathcal{H}_t = \mathcal{G}_t \vee \mathcal{F}_t, \text{ for all } t \in [0,T]. \tag{B.3.3}$$

B.3.1 Application to the Tandem System

In the context of the control problem in the main text, the construction above is specified in the following manner.

We temporarily fix the index n in the sequence of tandem stations whose queue lengths are modeled in (4.2.1) and (4.2.2) or (4.3.1). The exogenously driven processes are the arrival process into the first station and the potential service process from the second station. Therefore, the analogue of the expression (B.3.1) for the observation of the behavior of the uncontrolled processes is

$$\mathcal{G}^{(n)}_t = \sigma(N_1^+(n\mathcal{I}_s(\lambda)), N_2^-(n\mathcal{I}_s(\mu^2)) : s \leq t), \text{ for every } t \in [0,T]. \tag{B.3.4}$$

The only controllable random process is the potential service in the first station. Therefore, our construction requires only one copy of the family of independent processes indicated earlier. This simplifies the complete available information record (an analogue of (B.3.3)) to

$$\mathcal{H}_t^{(n)} = \mathcal{G}_t^{(n)} \vee \mathcal{F}_t^{(n)}, \text{ for every } t \in [0, T],$$

with

$$\mathcal{F}_t^{(n)} = \sigma(\xi(A) : A \in \mathcal{B}([0, t] \times [0, \infty))), \text{ for all } t \in [0, T],$$

and where ξ is a Poisson point process on $[0, t] \times [0, \infty)$, independent of N_1^+ and N_2^-, and with the Lebesgue measure as the intensity measure

Finally, let the set of all $\{\mathcal{H}_t^{(n)}\}-$predictable controls be denoted by $\mathcal{L}^{(n)}$. With the extra constraint m on the available cumulative service enforced in the main text, the space of all admissible controls is

$$\mathcal{L}^{(n)}(m) = \{\mu \in \mathcal{L}^{(n)} : \mathcal{I}_T(\mu) \leq m, \, a.s.\}.$$

To ease the notation of Section 4.8, we set

$$\mathcal{B} = \mathcal{L}^{(1)}(m). \tag{B.3.5}$$

Appendix C

The Single Station - Auxiliary Results

C.1 Monotonicity of Fluid-Regime Performance Measures

In this section we define the class of *monotone* performance measures in the fluid limit regime, as well as some properties pertaining to this class. This class will contain the performance measure specified in (3.2.5). The established properties of the whole class will facilitate the analysis of that particular performance measure.

C.1.1 Definitions

We shall be exploring a monotonicity feature on the space of performance measures. In order to facilitate the introduction of an ordering of possible queue lengths, let us define the following mappings.

Definition C.1.1. We define mappings $x, q : \mathbb{R}_+ \times \mathbb{L}^1_+[0,T] \times \mathbb{L}^1_+[0,T] \to \mathcal{C}$, as

$$x(q_0, \lambda, \mu) = q_0 + \mathcal{I}(\lambda - \mu) \text{ and } q = \Gamma \circ x,$$

for $q_0 \geq 0$ and λ and μ nonnegative integrable functions on the segment $[0,T]$. We say that the process $q : [0,T] \to \mathbb{R}$ is the fluid-limit queue length processes *generated* by λ as arrival rate and μ as service rate, starting at q_0.

In future notation, if there is no chance of confusion, we shall suppress the primitive functions describing the arrival and service rate, simply stating all the claims in terms of the resulting fluid limit queue.

Definition C.1.2. Let $q_1, q_2 \in \mathcal{C}$ be such that $q_1 \leq q_2$ almost everywhere. Then we say that q_2 *dominates* q_1, and write $q_1 \preceq q_2$.

Definition C.1.3. Let $\mathcal{D}om$ be a subset of $\mathbb{R}_+ \times \mathbb{L}^1_+[0,T] \times \mathbb{L}^1_+[0,T]$. The mapping $j : \mathcal{D}om \to \mathbb{R}$ is called an *increasing* (resp. *decreasing*) performance measure, if for all $(q_0^i, \lambda^i, \mu^i), i = 1, 2$, such that $q(q_0^1, \lambda^1, \mu^1) \preceq q(q_0^2, \lambda^2, \mu^2)$ (resp. $q(q_0^1, \lambda^1, \mu^1) \succeq q(q_0^2, \lambda^2, \mu^2)$) we have $j(q_0^1, \lambda^1, \mu^1) \leq j(q_0^2, \lambda^2, \mu^2)$ (resp. $j(q_0^1, \lambda^1, \mu^1) \geq j(q_0^2, \lambda^2, \mu^2)$). If j is either decreasing or increasing, it will be referred to as a *monotone* performance measure.

C.1.2 Examples

Let us take some time to describe a few legitimate (both in the formal and in the intuitive sense) performance measures which are either monotone or **not** monotone in the sense of Definition C.1.3.

Example C.1.4. [Holding cost] Consider a single station, single server system on a time-interval $[0,T]$. The arrival rate is denoted by λ and the service rate by μ. Initially, there are q_0 jobs at the station. Whenever there are jobs queued up at the station, a certain holding cost, depending on the number of queued up jobs, must be paid. Let us denote that cost by a function $h : \mathbb{R}_+ \to \mathbb{R}_+$. Obviously, only nondecreasing h make sense here.

Altogether, the aggregated cost over $[0,T]$ is given by

$$j(q_0, \lambda, \mu) = \int_0^T h(q_t(q_0, \lambda, \mu))\, dt.$$

Trivially, the performance measure j is monotone in the sense of Definition C.1.3.

Example C.1.5. [Average delay] Let us assume, for simplicity, that all the queues in this example are initially empty. The length of the queue generated by an arrival rate λ and a service rate μ can then be unambiguously denoted by $q(\lambda, \mu)$, through the obvious reduction of the notation introduced in Definition C.1.1. Furthermore, let us introduce the *delay* in $q(\lambda, \mu)$ at an instant $t \in [0,T]$ as

$$d_t(\lambda, \mu) = \inf\{\delta > 0 : \mathcal{I}_t(\lambda) = \mathcal{I}_{t-\delta}(\lambda) + q_t(\lambda, \mu)\}.$$

The *average delay* on the finite time-horizon $[0,T]$ is then given by

$$D(\lambda, \mu) = \frac{1}{T} \int_0^T d_t(\lambda, \mu)\, dt.$$

The average delay D as a real function of a pair in $\mathbb{L}^1_+[0,T] \times \mathbb{L}^1_+[0,T]$ is the performance measure we want to test for the monotonicity property of Definition C.1.3.

Consider arrival rates $\lambda_1 = \lambda_2 = \frac{1}{2}\lambda_3 \equiv 1$, and service rates $\mu_1 = 2\mu_2 = \frac{1}{2}\mu_3 \equiv \frac{1}{2}$. The queues generated by these rates are

$$q_1 = q(\lambda_1, \mu_1) = \frac{1}{2}\, e, \quad q_2 = q(\lambda_2, \mu_2) = \frac{3}{4}\, e, \quad \text{and} \quad q_3 = q(\lambda_3, \mu_3) = e.$$

Obviously, $q_1 \preceq q_2 \preceq q_3$. On the other hand, the delays for all three queues at any time t can easily be calculated and they equal

$$d_t(\lambda_1,\mu_1)=\frac{1}{2}\,t,\quad d_t(\lambda_2,\mu_2)=\frac{3}{4}\,t,\quad \text{and}\quad d_t(\lambda_3,\mu_3)=\frac{1}{2}\,t.$$

Therefore, we have $D(\lambda_1,\mu_1)\le D(\lambda_2,\mu_2)$ and $D(\lambda_3,\mu_3)\le D(\lambda_2,\mu_2)$. We conclude that average delay is not a monotone performance measure in the sense of Definition C.1.3.

Example C.1.6. [Lost goods and/or orders] We look at a manufacturing facility with a finite storage of size K, which is empty at the beginning of a production cycle $[0,T]$. The product is being manufactured at a rate λ and the rate of demand is given by μ.

Since the on-site storage is finite, all goods that are produced after the full capacity is reached are lost. In fact, the manufacturer needs to pay a penalty c per unit of lost goods (say, there is a disposal fee, or simply the extra product goes to waste and cannot make up for the cost of raw material used).

On the other hand, the manufacturer has an agreement with the merchant handling the product to pay a penalty d per unit of demanded goods that cannot be delivered due to a shortage in the storage. Also, there is no possibility of a backlog on the merchants orders - all orders that are not promptly completed are lost to the manufacturer (say, the merchant turns to another supplier).

We employ the expression for the two-sided regulator in (1.2.5) to display the amount of goods in the storage of the manufacturing facility as

$$\Gamma^K(x(\lambda,\mu))=x(\lambda,\mu)+l(\lambda,\mu)-u(\lambda,\mu),$$

where l and u are the regulator maps associated with $x(\lambda,\mu)$ and K as in Definition 1.2.2, and where we used obvious modifications of the notation introduced in Definition C.1.1.

Altogether, the manufacturer's cost on $[0,T]$, with rates λ and μ is

$$j(\lambda,\mu)=c\,u(\lambda,\mu)+d\,l(\lambda,\mu). \tag{C.1.1}$$

Let us check if j is a monotone performance measure.

Suppose that the storage capacity is $K=\frac{T}{4}$. Let $\lambda_1=\lambda_2=\lambda_3\equiv 1$, $\mu_1\equiv 1$, $\mu_2\equiv\frac{1}{2}$, and $\mu_3\equiv 2$. Then we have

$$q_1=q(\lambda_1,\mu_1)=0,\quad q_2=q(\lambda_2,\mu_2)=\frac{1}{2}\left(t\wedge\frac{T}{2}\right),\quad \text{and}\quad q_3=q(\lambda_3,\mu_3)=0.$$

Obviously, $q_1\preceq q_2$ and $q_3\preceq q_2$. The corresponding costs, as described in (C.1.1), are

$$j(\lambda_1,\mu_1)=0,\quad j(\lambda_2,\mu_2)=\frac{cT}{4},\quad \text{and}\quad j(\lambda_3,\mu_3)=\frac{d\,T}{2}.$$

For instance, the choice of parameters $c=1$ and $d=\frac{3}{4}$ produces a non-monotone performance measure.

Example C.1.7. Let the system consist of a single storage facility with the arrival rate of products equal to $\lambda \equiv 1$. The station supplies the demand stream μ. At the end of the production period $[0, T]$, if the remaining supplies exceed the level K_l, these have extra value when compared to the first K_l products. One can think of the first K_l products as dedicated to a costumer, and everything exceeding K_l as surplus that can produce extra profit. On the other hand, at all times at which the queue length exceeds $K_u > K_l$, an extra storage facility needs to be opened incurring a flat fee we denote by c. Summing the profit and cost we obtain the mapping $j : \mathbb{L}^1_+[0, T] \times \mathbb{L}^1_+[0, T] \to \mathbb{R}$, of the form

$$j(q) = (q_T(\lambda, \mu) - K_l)^+ - c \cdot \text{meas}\{t \in [0, T] : q_t(\lambda, \mu) > K_u\}. \tag{C.1.2}$$

Let us assume that $2K_u \leq T$ and consider three possible demand streams

$$\begin{aligned}
\mu_1 &\equiv 1, \\
\mu_2 &= 1 - \frac{K_u}{T}, \\
\mu_3(t) &= \left(1 - \frac{2K_u}{T}\right) \mathbf{1}_{[0, \frac{3T}{4}]}(t) + \left(1 + \frac{2K_u}{T}\right) \mathbf{1}_{(\frac{3T}{4}, T]}(t).
\end{aligned}$$

Along with a common arrival rate $\lambda \equiv 1$, these demand rates generate fluid-limit queue length processes

$$\begin{aligned}
q_1 &\equiv 0, \\
q_2(t) &= \frac{K_u}{T}\, t, \\
q_3(t) &= \frac{2K_u}{T}\, t \mathbf{1}_{[0, \frac{3T}{4}]}(t) + (3K_u - \frac{2K_u}{T}\, t) \mathbf{1}_{(\frac{3T}{4}, T]}(t).
\end{aligned}$$

Obviously, $q_1 \preceq q_2 \preceq q_3$, in the sense of Definition C.1.2. The mapping j defined in (C.1.2) evaluated at these queue length functions produces

$$\begin{aligned}
j(q_1) &= 0, \\
j(q_2) &= K_u - K_l, \\
j(q_3) &= K_u - K_l - \frac{T}{2}.
\end{aligned}$$

Hence, $j(q_1) < j(q_2)$ and $j(q_3) < j(q_2)$ and j is not monotone in the sense of Definition C.1.3.

C.1.3 Properties

We return to the performance measure considered in Section 3.2, as defined in (3.2.5). Let us fix an arrival rate λ and consider all initially empty fluid limit queues generated by varying the service rate μ across $\mathbb{L}^1_+[0, T]$. Formally, the mapping $Q^\lambda : \mathbb{L}^1_+[0, T] \to \mathcal{C}$ is introduced as

$$Q^\lambda(\mu) = \Gamma(\mathcal{I}(\lambda - \mu)), \text{ for } \mu \in \mathbb{L}^1_+[0, T].$$

We can also express Q^λ as $Q^\lambda(\mu) = q(0, \lambda, \mu)$, in the notation of Definition C.1.1.

Proposition C.1.8. *Let us denote by $\mathcal{D}om$ the set $\{0\} \times \{\lambda\} \times \mathbb{L}^1_+[0,T]$ and consider an increasing performance measure $j : \mathcal{D}om \to \mathbb{R}$. For a given constant m, define $\mu^* = \lambda \mathbf{1}_{[0,\tau(m)]}$, where τ is defined by (3.2.7). Then*

$$\inf_{\mu \in \bar{\mathcal{L}}(m)} j(Q^\lambda(\mu)) = j(Q^\lambda(\mu^*)),$$

where

$$\bar{\mathcal{L}}(m) = \{\mu \in \mathbb{L}^1_+[0,T] : \mathcal{I}(\mu)_T \leq m\}. \tag{C.1.3}$$

Proof. It suffices to prove that $Q^\lambda(\mu^*) \preceq Q^\lambda(\mu)$, for all $\mu \in \bar{\mathcal{L}}_m$. For all $t \leq \tau(m)$, $x_t(0, \lambda, \mu^*) = \mathcal{I}_t(\lambda - \mu^*) = 0$. Hence, for all $t \leq \tau(m)$, $Q^\lambda_t(\mu^*) = x_t(0, \lambda, \mu^*) = 0 \leq Q^\lambda_t(\mu)$.

On the other hand, for all $t \geq \tau(m)$, $\mathcal{I}_t(\mu^*) = \mathcal{I}_{\tau(m)}(\mu^*) = m$. Moreover, for all $\mu \in \bar{\mathcal{L}}(m)$, and all t, $\mathcal{I}_t(\mu) \leq m$, by the defining property (C.1.3). Combining these two inequalities, we get that for all $t > \tau(m)$, $Q^\lambda_t(\mu^*) = \mathcal{I}_t(\lambda) - m \leq \mathcal{I}_t(\lambda - \mu) = x_t(0, \lambda, \mu) \leq q_t(0, \lambda, \mu) = Q^\lambda_t(\mu)$.

The claim of the proposition follows from the assumed monotonicity of j. □

C.2 Inequalities. Tail Events.

C.2.1 General

An Inequality

Lemma C.2.1. *For all $a, b \in \mathbb{R}$, we have*

$$2^3(a^4 + b^4) \geq (a+b)^4 .$$

Proof. The function $f : \mathbb{R} \longrightarrow \mathbb{R}$, $f(x) = x^4$ is a convex function. Therefore, for all real numbers a, b

$$\frac{1}{2}(a^4 + b^4) \geq \left(\frac{a+b}{2}\right)^4 .$$

Rearranging the terms in this expression, we get the announced claim. □

On Properties Holding Eventually

The reader might wish to review the definition in (1.2.1).

Lemma C.2.2. *Let* $\{Y_n\}$ *and* $\{Z_n\}$ *be sequences of random elements on the probability space* $(\Omega, \mathcal{F}, \mathbb{P})$ *taking values in a measurable space* $(S, \mathcal{S})$*. Suppose that* $\{Y_n\}$ *and* $\{Z_n\}$ *satisfy*

$$\mathbb{P}[Y_n = Z_n,\, ev.] = 1. \tag{C.2.1}$$

Let A *be any measurable subset of* S*, then the following claim holds*

$$\lim_{n\to\infty} [\mathbb{P}[Y_n \in A] - \mathbb{P}[Z_n \in A]] = 0. \tag{C.2.2}$$

Moreover, we have that

$$\lim_{n\to\infty} \mathbb{P}[Y_n \in A] = \lim_{n\to\infty} \mathbb{P}[Z_n \in A], \tag{C.2.3}$$

in the sense that if any of the above limits exists, then they both exist and are equal.

Proof. For any n, we have

$$\mathbb{P}[Y_n \in A] = \mathbb{E}[\mathbf{1}_{[Y_n\in A]}] = \mathbb{E}[\mathbf{1}_{[Y_n\in A]}\mathbf{1}_{[Y_n=Z_n]}] + \mathbb{E}[\mathbf{1}_{[Y_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}]. \tag{C.2.4}$$

Similarly,

$$\mathbb{P}[Z_n \in A] = \mathbb{E}[\mathbf{1}_{[Z_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}] + \mathbb{E}[\mathbf{1}_{[Z_n\in A]}\mathbf{1}_{[Y_n=Z_n]}]. \tag{C.2.5}$$

Therefore, the absolute value of the difference of probabilities in (C.2.2) equals

$$\begin{aligned}
&|\mathbb{P}[Y_n \in A] - \mathbb{P}[Z_n \in A]| \\
&\qquad = |\mathbb{E}[\mathbf{1}_{[Y_n\in A]}\mathbf{1}_{[Y_n=Z_n]}] + \mathbb{E}[\mathbf{1}_{[Y_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}] \\
&\qquad\qquad - (\mathbb{E}[\mathbf{1}_{[Z_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}] + \mathbb{E}[\mathbf{1}_{[Z_n\in A]}\mathbf{1}_{[Y_n=Z_n]}])| \\
&\qquad = |\mathbb{E}[\mathbf{1}_{[Y_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}] - \mathbb{E}[\mathbf{1}_{[Z_n\in A]}\mathbf{1}_{[Y_n\neq Z_n]}]| \\
&\qquad \leq 2\mathbb{E}[\mathbf{1}_{[Y_n\neq Z_n]}] \\
&\qquad = 2\mathbb{P}[Y_n \neq Z_n].
\end{aligned}$$

By assumption (C.2.1), we have

$$\mathbb{E}[\mathbf{1}_{[Y_n\neq Z_n]}] \to 0, \text{ as } n \to \infty. \tag{C.2.6}$$

The last two displays imply claim (C.2.2).

Putting together this claim with the existence of one of the limits from (C.2.3) completes the proof of the lemma. □

C.2.2 Inequalities involving Brownian Motion

Lemma C.2.3. *Let* $(B_t^{(n)})_{t\in\mathbb{R}_+}$ *be a sequence of standard Brownian motions on a common probability space. Let* $\{\alpha_n\}$ *be a sequence in* $[0,T]$ *and* $\{f_n\}$ *be a sequence of nondecreasing, deterministic functions* $f_n : [0,T] \to \mathbb{R}$ *such that* $f_n(\alpha_n) > 0$*, and* $f_n(T) \le C$*, for some positive constant* C*. Moreover, let* $\{s_n\}$ *and* $\{c_n\}$ *be sequences of positive real numbers satisfying* $\sum_{n=1}^{\infty} \frac{c_n}{s_n} e^{-\frac{s_n^2}{2C}} < \infty$*. Then we have, as* $n \to \infty$*,*

$$S^{(n)} := c_n meas\{t \in [\alpha_n, T] : B^{(n)}_{f_n(t)} > s_n\} \to 0,\ a.s. \tag{C.2.7}$$

Proof. By the Borel-Cantelli lemma, it suffices to prove that for any given $\epsilon > 0$

$$\sum_{n=1}^{\infty} \mathbb{P}[S^{(n)} > \epsilon] < \infty.$$

For any fixed n, the probability in question reads as

$$\mathbb{P}[S^{(n)} > \epsilon] = \mathbb{P}\left[\text{meas}\{t \in [\alpha_n, T] : B^{(n)}_{f_n(t)} > s_n\} > \frac{\epsilon}{c_n}\right].$$

By Markov's inequality this quantity is dominated by

$$\frac{c_n}{\epsilon}\,\mathbb{E}[\text{meas}\{t \in [\alpha_n, T] : B^{(n)}_{f_n(t)} > s_n\}] = \frac{c_n}{\epsilon}\int_{\alpha_n}^{T} \mathbb{P}[B^{(n)}_{f_n(t)} > s_n]dt, \tag{C.2.8}$$

where the last equality holds by Fubini's theorem. Let us consider a single instance $t \in [\alpha_n, T]$. We have

$$\mathbb{P}\left[B^{(n)}_{f_n(t)} > s_n\right] = \mathbb{P}\left[\frac{B^{(n)}_{f_n(t)}}{\sqrt{f_n(t)}} > \frac{s_n}{\sqrt{f_n(t)}}\right] = \mathbb{P}\left[Y > \frac{s_n}{\sqrt{f_n(t)}}\right],$$

where Y is a standard normal random variable. Furthermore, using the fact that f_n is nondecreasing and bounded by C, we obtain the following upper bound on the above probability, which is independent of t:

$$\mathbb{P}\left[Y > \frac{s_n}{\sqrt{C}}\right].$$

Returning to equation (C.2.8), we obtain the following upper bound on its right-hand side:

$$\frac{c_n T}{\epsilon}\,\mathbb{P}\left[Y > \frac{s_n}{\sqrt{C}}\right] = \frac{c_n T}{\epsilon}\int_{\frac{s_n}{\sqrt{C}}}^{\infty} \varphi(y)dy,$$

where φ represents the density of the standard normal distribution. Since s_n are assumed to be positive, this expression is further dominated by

$$\frac{c_n T}{\epsilon} \frac{\sqrt{C}}{s_n} \varphi\left(\frac{s_n}{\sqrt{C}}\right) = \frac{T\sqrt{C}}{\epsilon\sqrt{2\pi}} \frac{c_n}{s_n} e^{-\frac{s_n^2}{2C}}.$$

Summing up over $n \in \mathbb{N}$, and using the assumption on the sequences $\{s_n\}$ and $\{c_n\}$ we obtain the desired inequality

$$\sum_{n=1}^{\infty} \mathbb{P}[S^{(n)} > \epsilon] \leq \frac{T\sqrt{C}}{\epsilon\sqrt{2\pi}} \sum_{n=1}^{\infty} \frac{c_n}{s_n} e^{-\frac{s_n^2}{2C}} < \infty.$$

□

Corollary C.2.4. *Let $(B_t^{(n)})_{t\in\mathbb{R}_+}$ be a sequence of standard Brownian motions on a common probability space. Let $\{\alpha_n\}$ be a sequence in $[0,T]$ and $\{f_n\}$ be a sequence of nondecreasing, deterministic functions $f_n : [0,T] \to \mathbb{R}$ such that $f_n(\alpha_n) > 0$ and $f_n(T) \leq C$, for some positive constant C. Let $\{s_n\}$ be a sequence of positive real numbers satisfying $\sum\limits_{n=1}^{\infty} \frac{1}{s_n} e^{-\frac{s_n^2}{2C}} < \infty$. Then we have, as $n \to \infty$,*

$$S^{(n)} := meas\{t \in [\alpha_n, T] : B_{f_n(t)}^{(n)} > s_n\} \to 0, \ a.s. \tag{C.2.9}$$

Lemma C.2.5. *Let $(B_t)_{t\in\mathbb{R}_+}$ be a standard Brownian motion. Let $\{\alpha_n\}$ be a sequence in $[0,T]$ and $\{f_n\}$ be a sequence of nondecreasing, deterministic functions $f_n : [0,T] \to \mathbb{R}$ such that $f_n(\alpha_n) > 0$ and $f_n(T) \leq C$, for some positive constant C. Moreover, let $\{s_n\}$ be a divergent sequence of positive real numbers. Then, as $n \to \infty$,*

$$S^{(n)} := meas\{t \in [\alpha_n, T] : B_{f_n(t)} > s_n\} \to 0,$$

in expectation.

Proof. By Fubini's theorem, it follows that

$$\mathbb{E}[S^{(n)}] = \mathbb{E}[\text{meas}\{t \in [\alpha_n, T] : B_{f_n(t)} > s_n\}] = \int_{\alpha_n}^{T} \mathbb{P}[B_{f_n(t)} > s_n] dt. \tag{C.2.10}$$

Considering any $t \in [\alpha_n, T]$, we observe that

$$\mathbb{P}[B_{f_n(t)} > s_n] = \mathbb{P}\left[\frac{B_{f_n(t)}}{\sqrt{f_n(t)}} > \frac{s_n}{\sqrt{f_n(t)}}\right] = \mathbb{P}\left[Y > \frac{s_n}{\sqrt{f_n(t)}}\right],$$

where Y is a standard normal random variable. Recalling that f_n is nondecreasing and bounded by C, an upper bound on the above probability is

$$\mathbb{P}\left[Y > \frac{s_n}{\sqrt{C}}\right].$$

We can now dominate the right-hand side of (C.2.10) by the quantity

$$T\,\mathbb{P}\left[Y > \frac{s_n}{\sqrt{C}}\right] = T \int_{\frac{s_n}{\sqrt{C}}}^{\infty} \varphi(y)dy,$$

where φ represents the density of the standard normal distribution. Since s_n are assumed to be positive, this expression is further dominated by

$$\frac{T\sqrt{C}}{s_n}\varphi\left(\frac{s_n}{\sqrt{C}}\right) = \frac{T\sqrt{C}}{\sqrt{2\pi}}\frac{1}{s_n}e^{-\frac{s_n^2}{2C}}.$$

As the sequence $\{s_n\}$ is assumed to be divergent, the proof is finished. □

We now wish to be able to substitute the term s_n in (C.2.9) by a term that depends on t as well. For simplicity, we discard the dependence of the "time-change" on n.

Lemma C.2.6. *Let $(B_t)_{t\in\mathbb{R}_+}$ be a Brownian motion and $\{\alpha_n\}$ a sequence in $[0,T]$. Let $f : [0,T] \to \mathbb{R}$ be a nondecreasing, deterministic function such that $f(\alpha_n) > 0$, for all n, and let $g : [0,T] \to \mathbb{R}$ be a nonincreasing, deterministic function such that $g(t) > 0$, for all $t < T$. Finally, let $\{s_n\}$ be a divergent sequence of positive real numbers. Then, as $n \to \infty$,*

$$S^{(n)} := \mathit{meas}\{t \in [\alpha_n, T] : B_{f(t)} > s_n g(t)\} \to 0 \tag{C.2.11}$$

in expectation.

Proof. In case that $g(T) > 0$, the random variable $S^{(n)}$ is dominated by

$$S^{(n)} = \text{meas}\{t \in [\alpha_n, T] : B_{f(t)} > s_n g(T)\},$$

which brings us back to the realm of Lemma C.2.5.

The case when $g(T) = 0$ requires a bit more work. Let $\epsilon > 0$ be an arbitrary constant, and let $\delta = g(T-\epsilon) > 0$. $S^{(n)}$ can now be conveniently rewritten as

$$\begin{aligned} S^{(n)} &= \text{meas}\{t \in [\alpha_n, T-\epsilon) : B_{f(t)} > s_n g(t)\} + \text{meas}\{t \in (T-\epsilon, T] : B_{f(t)} > s_n g(t)\} \\ &\le \text{meas}\{t \in [\alpha_n, T-\epsilon) : B_{f(t)} > s_n \delta\} + \epsilon. \end{aligned}$$

An application of Lemma C.2.5 and the arbitrariness of ϵ complete the proof. □

C.2.3 Inequalities Involving Poisson Processes

Lemma C.2.7. *Let Y be a unit Poisson process. Then for all positive constants ξ,*

$$\mathbb{E}\left[(Y_\xi - \xi)^4\right] = 3\xi^2 + \xi.$$

Proof. Using the facts that the excess kurtosis of the Poisson distribution with parameter ξ equals $1/\xi$, and that its variance is ξ, we get

$$\mathbb{E}\left[(Y_\xi - \xi)^4\right] = \xi^2(3 + \frac{1}{\xi}) = 3\xi^2 + \xi.$$

□

Lemma C.2.8. *Let $Y^{(1)}$ and $Y^{(2)}$ be independent nonhomogeneous Poisson processes, with means $y^{(1)}$ and $y^{(2)}$. Assume that $0 < \alpha < \beta$ are real constants such that*

$$\inf_{\alpha<t<\beta} \mathcal{I}_t(y^{(1)} - y^{(2)}) =: \underline{y} > 0. \tag{C.2.12}$$

Then the following inequality holds true

$$\mathbb{P}[\inf_{\alpha<t<\beta}[Y_t^{(1)} - Y_t^{(2)}] \leq 0] \leq \frac{8}{\underline{y}^4}[3(\mathcal{I}_\beta(y^{(1)})^2 + \mathcal{I}_\beta(y^{(2)})^2) + \mathcal{I}_\beta(y^{(1)} + y^{(2)})].$$

Proof. Let us first compensate the Poisson processes on the left-hand side of (C.2.13), thus obtaining a martingale, and use condition (C.2.12) to obtain an upper bound. Indeed, we have

$$\begin{aligned}
&\mathbb{P}[\inf_{\alpha<t<\beta}[Y_t^{(1)} - \mathcal{I}_t(y^{(1)}) - (Y_t^{(2)} - \mathcal{I}_t(y^{(2)})) + \mathcal{I}_t(y^{(1)} - y^{(2)})] \leq 0] \\
&\qquad \leq \mathbb{P}\Big[\inf_{\alpha<t<\beta}[Y_t^{(1)} - \mathcal{I}_t(y^{(1)}) - (Y_t^{(2)} - \mathcal{I}_t(y^{(2)}))] \leq -\underline{y}\Big] \\
&\qquad \leq \mathbb{P}\Big[\sup_{\alpha<t<\beta}[Y_t^{(1)} - \mathcal{I}_t(y^{(1)}) - (Y_t^{(2)} - \mathcal{I}_t(y^{(2)}))]^4 \geq \underline{y}^4\Big].
\end{aligned} \tag{C.2.13}$$

Using the submartingale inequality, a further upper bound for the above expression is

$$\begin{aligned}
&\mathbb{P}\Big[\sup_{\alpha<t<\beta}[Y_t^{(1)} - \mathcal{I}_t(y^{(1)}) - (Y_t^{(2)} - \mathcal{I}_t(y^{(2)}))]^4 \geq \underline{y}^4\Big] \\
&\qquad \leq \frac{1}{\underline{y}^4}\mathbb{E}[(Y_\beta^{(1)} - \mathcal{I}_\beta(y^{(1)}) - (Y_\beta^{(2)} - \mathcal{I}_\beta(y^{(2)})))^4].
\end{aligned}$$

Using Lemma C.2.7, we see that

$$\begin{aligned}
&\mathbb{E}[(Y_\beta^{(1)} - \mathcal{I}_\beta(y^{(1)}) - (Y_\beta^{(2)} - \mathcal{I}_\beta(y^{(2)})))^4] \\
&\quad \leq 8[\mathbb{E}[(Y_\beta^{(1)} - \mathcal{I}_\beta(y^{(1)}))^4] + \mathbb{E}[(Y_\beta^{(2)} - \mathcal{I}_\beta(y^{(2)}))^4]] \\
&\quad = 8[3(\mathcal{I}_\beta(y^{(1)})^2 + \mathcal{I}_\beta(y^{(2)})^2) + \mathcal{I}_\beta(y^{(1)} + y^{(2)})].
\end{aligned}$$

The lemma is a direct consequence of the last three displays. □

C.2.4 Inequalities Involving Submartingales

Lemma C.2.9. *Let* $(S_t)_{t\in\mathbb{R}_+}$ *be a submartingale. Let* $\{\alpha_n\}$ *be a sequence in* $[0,T]$*, and let* $\{f_n\}$ *be a sequence of nondecreasing, deterministic functions* $f_n : [0,T] \to \mathbb{R}$ *such that* $f_n(\alpha_n) > 0$*, and* $f_n(T) \leq C$ *for some positive constant* C*. Finally, let* $\{s_n\}$ *and* $\{c_n\}$ *be sequences of positive real numbers satisfying* $\sum_{n=1}^{\infty} \frac{c_n}{s_n} < \infty$*. Then we have*

$$S^{(n)} := c_n meas\{t \in [\alpha_n, T] : S_{f_n(t)} > s_n\} \to 0, \ a.s.$$

Proof. It suffices to prove that for all $\epsilon > 0$,

$$\mathbb{P}[S^{(n)} > \epsilon, \text{ i.o.}] = 0.$$

By the Borel-Cantelli Lemma this is equivalent to the inequality

$$\sum_{n=1}^{\infty} \mathbb{P}[S^{(n)} > \epsilon] < \infty.$$

Temporarily fixing $\epsilon > 0$ and $n \in \mathbb{N}$, and using the Markov inequality followed by Fubini's theorem, we get

$$\mathbb{P}[S^{(n)} > \epsilon] \leq \frac{\mathbb{E}[S^{(n)}]}{\epsilon} = \frac{c_n}{\epsilon} \int_{\alpha_n}^{T} \mathbb{P}[S_{f_n(t)} > s_n]\, dt.$$

Next, we perform the following short calculation finalized by the submartingale inequality

$$\begin{aligned}
\mathbb{P}[S^{(n)} > \epsilon] &\leq \frac{c_n(T-\alpha_n)}{\epsilon} \mathbb{P}\left[\sup_{t\in[\alpha_n,T]} S_{f_n(t)} > s_n\right] \\
&\leq \frac{c_n T}{\epsilon} \mathbb{P}\left[\sup_{u\in[f_n(\alpha_n),\, f_n(T)]} S_u > s_n\right] \\
&\leq \frac{c_n T}{\epsilon} \mathbb{P}\left[\sup_{u\in[0,C]} S_u > s_n\right] \\
&\leq \frac{c_n T \mathbb{E}[S_C^+]}{\epsilon s_n}.
\end{aligned}$$

Summing over all n, we get

$$\sum_{n=1}^{\infty} \mathbb{P}[S^{(n)} > \epsilon] \leq \frac{T\,\mathbb{E}[S_C^+]}{\epsilon} \sum_{n=1}^{\infty} \frac{c_n}{s_n} < \infty.$$

□

C.3 Limiting Lower Bound - Second Order Analysis

Throughout this section W denotes a standard Brownian motion on the segment $[0, T]$ and the standard normal distribution function is denoted by F_N. Also, the standard notation of uniform acceleration (see [Whi02a], e.g.) is used without explicit reiteration of definitions.

In the present section we commit our attention to the infinite-buffer, single-station model studied in the main text. After the fluid-limit queue length process leaves zero for the last time in the segment $[0, T]$, the fluctuation process in the second order approximation is more likely to be positive than negative. We dedicate the first lemma in this section to this useful observation.

Lemma C.3.1. *Let $\mu \in \mathbb{L}^1_+[0, T]$ be a deterministic service rate, and let $t^* = \sup\{t \in [0, T] : \bar{Q}_t(\mu) = 0\} \wedge T$. Then for all $\kappa \geq 0$ and t, such that $T > t \geq t^*$,*

$$\mathbb{P}[\hat{Q}_t(\mu) > \kappa] \geq 1 - F_N\left(\frac{\kappa}{\sqrt{\Sigma_t - \Sigma_{t^*}}}\right),$$

where $\Sigma = \mathcal{I}(\lambda + \mu)$.

Proof. Since for all $t > t^*$, $\bar{Q}_t(\mu) > 0$, we conclude that

$$\Phi_{-\bar{X}(\mu)}(t) = \Phi_{-\bar{X}(\mu)}(t^*).$$

Let $\bar{t} = \sup \Phi_{-\bar{X}(\mu)}(t^*)$. The set $\Phi_{-\bar{X}(\mu)}(t^*)$ is compact (see Corollary 9.3.1. in [Whi02a]), so it contains its supremum $\bar{t}$. Along with the defining expression (3.3.3), this allows us to write

$$\begin{aligned}\hat{Q}_t(\mu) &= W(\mathcal{I}_t(\lambda) + \mathcal{I}_t(\mu)) - W(\mathcal{I}_{\bar{t}}(\lambda) + \mathcal{I}_{\bar{t}}(\mu)) \\ &\quad + W(\mathcal{I}_{\bar{t}}(\lambda) + \mathcal{I}_{\bar{t}}(\mu)) + \sup_{s \in \Phi_{-\bar{X}(\mu)}(t)} [-W(\mathcal{I}_s(\lambda) + \mathcal{I}_s(\mu))], \\ &\geq W(\mathcal{I}_t(\lambda) + \mathcal{I}_t(\mu)) - W(\mathcal{I}_{\bar{t}}(\lambda) + \mathcal{I}_{\bar{t}}(\mu)),\end{aligned} \tag{C.3.1}$$

for all $t \geq t^*$. For any such fixed t, the random variable on the right-hand side of (C.3.1) is an increment of a Brownian motion. Hence, it is normally distributed, centered at zero with variance $\Sigma_t - \Sigma_{\bar{t}} = \mathcal{I}_t(\lambda + \mu) - \mathcal{I}_{\bar{t}}(\lambda + \mu)$. Therefore, for any $\kappa \geq 0$

$$\begin{aligned}\mathbb{P}[\hat{Q}_t(\mu) > \kappa] &\geq \mathbb{P}\left[Y > \frac{\kappa}{\sqrt{\Sigma_t - \Sigma_{\bar{t}}}}\right] \geq \mathbb{P}\left[Y > \frac{\kappa}{\sqrt{\Sigma_t - \Sigma_{t^*}}}\right] \\ &= 1 - F_N\left(\frac{\kappa}{\sqrt{\Sigma_t - \Sigma_{t^*}}}\right),\end{aligned}$$

where Y denotes a standard normal random variable. □

Remark C.3.1. One is tempted to try to work through the simple argument above in the case of nondeterministic μ. However, this plan requires that the nondeterministic analogue of $\bar{t}$ be a stopping time. Otherwise, there is no way to gain information on the distribution of the random variable on the right-hand side of C.3.1. For a generic nondeterministic μ, there is no reason that this should be true.

We continue by considering the behavior of the system in the two regions that are "beyond our control", i.e., beyond time $\tau((K+m)-)$, when the fluid limit is inevitably at the threshold or above it.

Lemma C.3.2. *Let $\{\mu_n\}$ be a deterministic admissible sequence. Then necessarily*

$$\liminf_{n\to\infty} \mathbb{E}\left[meas\left\{t\in[0,T] : \mathcal{I}_t(\lambda) = K+m, \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\right\}\right] \geq \frac{1}{2}(\tau(K+m) - \tau((K+m)-)).$$

Proof. Let $\{\mu_n\}$ be an arbitrary deterministic admissible sequence. For all n, we can employ Fubini's theorem to rewrite the amount of time spent above the threshold K on the set $\{t : \mathcal{I}_t(\lambda) = K+m\}$ as

$$\begin{aligned}\mathbb{E}\Big[\int_{\tau((K+m)-)}^{\tau(K+m)} \mathbf{1}_{\{\bar{Q}_t(\mu_n)+\frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n)>K\}}\,dt\Big] \\ = \int_{\tau((K+m)-)}^{\tau(K+m)} \mathbb{P}\left[\bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\right] dt.\end{aligned} \tag{C.3.2}$$

Noticing that for all $t \geq \tau((K+m)-)$ we have

$$\bar{Q}_t(\mu_n) \geq \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu_n) \geq K+m-m = K, \tag{C.3.3}$$

we realize that the value in (C.3.2) dominates the quantity

$$\int_{\tau((K+m)-)}^{\tau(K+m)} \mathbb{P}[\hat{Q}_t(\mu_n) > 0]\,dt. \tag{C.3.4}$$

Moreover, the inequality in (C.3.3) implies strict positivity of $\bar{Q}_t(\mu_n)$, for all $t \geq \tau((K+m)-)$. This validates the employment of Lemma C.3.1, which combined with (C.3.4) gives us the final lower bound on (C.3.2)

$$\frac{1}{2}\int_{\tau((K+m)-)}^{\tau(K+m)} dt = \frac{1}{2}(\tau(K+m) - \tau((K+m)-)).$$

□

Just following the introduction of performance measures $J^{(n)}$, we interpreted them as aggregated unit penalty incurred over the period when the queue length exceeds a given threshold. Using this nomenclature, we can say that Lemma C.3.2 provides a lower bound for the penalty accrued before time $\tau(K+m)$.

The next lemma is dedicated to the very end of the segment $[0,T]$. As it turns out, in the limit, the system will almost surely be paying full penalty in this region, regardless of the service discipline employed, as it did in the fluid optimization problem.

Lemma C.3.3. *For all deterministic admissible sequences* $\{\mu_n\}$,

$$meas\left\{t \in [\tau(K+m),T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\right\} \to T - \tau(K+m), \tag{C.3.5}$$

almost surely, as $n \to \infty$.

Proof. The fact that the left-hand side of (C.3.5) is bounded from above by $T - \tau(K+m)$ leads us to the conclude that it suffices to prove

$$\begin{aligned}\liminf_{n\to\infty} \mathbb{E}[\text{meas}\{t \in [\tau(K+m),T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\}] \\ \geq T - \tau(K+m),\end{aligned} \tag{C.3.6}$$

in order to prove the full statement of the lemma.

Focusing now on the proof of (C.3.6), we immediately see that, in the case of $T = \tau(K+m)$, the claim trivially holds true. Hence, we concentrate on the case of $\tau(K+m) < T$.

Let $\{\mu_n\}$ be an arbitrary sequence of admissible disciplines. Then, for all $t > \tau(K+m)$ and all $n \in \mathbb{N}$,

$$\bar{Q}_t(\mu_n) \geq \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu_n) > K + m - m = K.$$

Let us temporarily fix a sufficiently small, positive number ϵ, satisfying $\tau(K+m) + \epsilon \leq T$. We define $\delta := \mathcal{I}_{\tau(K+m)+\epsilon}(\lambda) - (K+m)$, and conclude straight from the definition of the mapping τ that δ is strictly positive. Then for all $t \geq \tau(K+m) + \epsilon$ and all n

$$\bar{Q}_t(\mu_n) \geq \mathcal{I}_t(\lambda) - \mathcal{I}_t(\mu_n) > K + m + \delta - m = K + \delta.$$

Hence, for all n,

$$\begin{aligned}&\text{meas}\left\{t \in [\tau(K+m),T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\right\} \\ &\geq \text{meas}\left\{t \in [\tau(K+m)+\epsilon,T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K\right\} \\ &\geq \text{meas}\left\{t \in [\tau(K+m)+\epsilon,T] : \hat{Q}_t(\mu_n) > -\delta\sqrt{n}\right\}.\end{aligned} \tag{C.3.7}$$

Let us temporarily fix an $n \in \mathbb{N}$. Since the fluid limit is in overload, $\bar{L}_t(\mu_n) = \bar{L}_T(\mu_n)$, for all $t \geq \tau(K+m)+\epsilon$. Hence, for all such t,

$$\Phi_{-\bar{X}(\mu_n)}(t) = \Phi_{-\bar{X}(\mu_n)}(T).$$

Let $t^*(\mu_n) = \sup \Phi_{-\bar{X}(\mu_n)}(T)$. By Lemma 9.3.3 in [Whi02a], $\Phi_{-\bar{X}(\mu)}(T)$ is a compact set, so $t^*(\mu_n) \in \Phi_{-\bar{X}(\mu)}(T)$. This allows us to rewrite $\hat{Q}_t(\mu_n)$, for $t \geq \tau(K+m)+\epsilon$, as

$$\begin{aligned}
\hat{Q}_t(\mu_n) &= W(\mathcal{I}_t(\lambda)+\mathcal{I}_t(\mu_n)) + \sup_{s\in\Phi_{-\bar{X}(\mu_n)}(T)} (-W(\mathcal{I}_s(\lambda)+\mathcal{I}_s(\mu_n))) \\
&= W(\mathcal{I}_t(\lambda)+\mathcal{I}_t(\mu_n)) - W(\mathcal{I}_{t^*(\mu_n)}(\lambda)+\mathcal{I}_{t^*(\mu_n)}(\mu_n)) \\
&\qquad + W(\mathcal{I}_{t^*(\mu_n)}(\lambda)+\mathcal{I}_{t^*(\mu_n)}(\mu_n)) + \sup_{s\in\Phi_{-\bar{X}(\mu_n)}(T)} (-W(\mathcal{I}_s(\lambda)+\mathcal{I}_s(\mu_n))) \\
&\geq W(\mathcal{I}_t(\lambda)+\mathcal{I}_t(\mu_n)) - W(\mathcal{I}_{t^*(\mu_n)}(\lambda)+\mathcal{I}_{t^*(\mu_n)}(\mu_n)).
\end{aligned} \tag{C.3.8}$$

In order to make the notation less cumbersome, let us introduce the process

$$\hat{W}^{(n)}_t = W(\mathcal{I}_{t+t^*(\mu_n)}(\lambda)+\mathcal{I}_{t+t^*(\mu_n)}(\mu_n)) - W(\mathcal{I}_{t^*(\mu_n)}(\lambda)+\mathcal{I}_{t^*(\mu_n)}(\mu_n)),$$

for $t \in [0, T-t^*(\mu_n)]$. This process depends on μ_n, but let us suppress that dependence from the notation for now. Evidently, $\hat{W}^{(n)}$ is a Brownian motion with the variance process

$$\Sigma^{(n)}_t = \mathcal{I}_{t+t^*(\mu_n)}(\lambda)+\mathcal{I}_{t+t^*(\mu_n)}(\mu_n) - (\mathcal{I}_{t^*(\mu_n)}(\lambda)+\mathcal{I}_{t^*(\mu_n)}(\mu_n)).$$

Finally, the estimate in (C.3.8) produces the lower bound of the form

$$\begin{aligned}
&\operatorname{meas}\{t\in[\tau(K+m)+\epsilon, T] : \hat{Q}_t(\mu_n) > -\delta\sqrt{n}\} \\
&\quad \geq \operatorname{meas}\{t\in[\tau(K+m)+\epsilon-t^*(\mu_n),\, T-t^*(\mu_n)] : \hat{W}^{(n)}_t > -\delta\sqrt{n}\} \\
&\quad = T-(\tau(K+m)+\epsilon) \\
&\qquad - \operatorname{meas}\{t\in[\tau(K+m)+\epsilon-t^*(\mu_n),\, T-t^*(\mu_n)] : \hat{W}^{(n)}_t \leq -\delta\sqrt{n}\} \\
&\quad = T-(\tau(K+m)+\epsilon) \\
&\qquad - \operatorname{meas}\{t\in[\tau(K+m)+\epsilon-t^*(\mu_n),\, T-t^*(\mu_n)] : -\hat{W}^{(n)}_t \geq \delta\sqrt{n}\} \\
&\quad = T-(\tau(K+m)+\epsilon) \\
&\qquad - \operatorname{meas}\left\{t\in[\tau(K+m)+\epsilon-t^*(\mu_n),\, T-t^*(\mu_n)] : -\frac{\hat{W}^{(n)}_t}{\sqrt{\Sigma^{(n)}_t}} \geq \frac{\delta\sqrt{n}}{\sqrt{\Sigma^{(n)}_t}}\right\}.
\end{aligned} \tag{C.3.9}$$

Note that for all $t \in [\tau(K+m)+\epsilon-t^*(\mu_n),\, T-t^*(\mu_n)]$, we have

$$\Sigma^{(n)}_t > K+m+\mathcal{I}_{t+t^*(\mu_n)}(\mu_n) - \mathcal{I}_{t^*(\mu_n)}(\lambda-\mu_n) - 2\mathcal{I}_{t^*(\mu_n)}(\mu_n). \tag{C.3.10}$$

Since $\bar{Q}_{t^*(\mu_n)}(\mu_n) = 0$ by definition of t^*, and $\bar{Q}_{t^*(\mu_n)}(\mu_n) \geq \mathcal{I}_{t^*(\mu_n)}(\lambda - \mu_n)$ due to the structure of the Skorokhod map, we conclude that (C.3.9) implies

$$\Sigma_t^{(n)} > K + m + \mathcal{I}_{t+t^*(\mu_n)}(\mu_n) - 2\mathcal{I}_{t^*(\mu_n)}(\mu_n) \geq K + m - \mathcal{I}_{t^*(\mu_n)}(\mu_n). \tag{C.3.11}$$

The last inequality and the fact that all μ_n are assumed to be admissible, and hence conform to the upper bound m on the total amount of available service, give us $\Sigma_t^{(n)} > K$. Therefore, all the divisions by $\Sigma_t^{(n)}$ in (C.3.9) are allowed. On the other hand, since $\Sigma_t^{(n)} \leq \mathcal{I}_T(\lambda) + m$, we get the following lower bound on the value in (C.3.11):

$$\begin{aligned} &T - (\tau(K+m) + \epsilon) \\ &\quad - \text{meas}\left\{ t \in [\tau(K+m) + \epsilon - t^*(\mu_n),\, T - t^*(\mu_n)] : -\frac{\hat{W}_t^{(n)}}{\sqrt{\Sigma_t^{(n)}}} \geq \frac{\delta\sqrt{n}}{\sqrt{\mathcal{I}_T(\lambda) + m}} \right\}. \end{aligned}$$

We now invoke the notation and results of Corollary C.2.4, setting $s_n = \delta\sqrt{n}/\sqrt{\mathcal{I}_T(\lambda) + m}$ and f_n to be the identity function for all n, and realizing that $-\hat{W}^{(n)}/\Sigma^{(n)}$ is a standard Brownian motion. Due to the arbitrary choice of ϵ, the proof is finished since all the assumptions of Corollary C.2.4 can be easily verified. □

Albeit interesting in its own right, the previous lemma is stronger than what is needed at present. The following corollary contains the exact result that will be referred to in the future.

Corollary C.3.4. *For all admissible sequences $\{\mu_n\}$, as $n \to \infty$,*

$$\mathbb{E}\left[meas\left\{ t \in [\tau(K+m), T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K \right\} \right] \to T - \tau(K+m).$$

Proof. The claim follows immediately from Lemma C.3.3 and from boundedness between 0 and T of random variables representing time spent above the threshold

$$\text{meas}\left\{ t \in [\tau(K+m), T] : \bar{Q}_t(\mu_n) + \frac{1}{\sqrt{n}}\hat{Q}_t(\mu_n) > K \right\}.$$

□

Finally, we prove that, in the limit, the expected penalty incurred by the sequence $\{\hat{\mu}_n\}$ matches the expectation of $\hat{J}^*$.

Proposition C.3.5. *Let $\hat{J}^*$ be as in (3.3.6). Then for the sequence $\{\hat{\mu}_n\}$ defined by (3.3.7)*

$$\lim_{n\to\infty} \mathbb{E}[\hat{J}^{(n)}(\hat{\mu}_n) - \hat{J}^*] = 0.$$

Proof. For typographical reasons, let us set $\tau_- = \tau((K+m)-)$ and $\tau_+ = \tau(K+m)$. Using (3.3.8) and (3.3.9), we see that the performance measure sequence $\{\hat{J}^{(n)}\}$, evaluated at $\{\hat{\mu}_n\}$, can be rewritten as

$$\begin{aligned}\hat{J}^{(n)}(\hat{\mu}_n) = \text{meas}&\left\{t \in [0, \tau(\eta_n K)) : W(\mathcal{I}_t(\lambda)) > \sqrt{n}(K - \mathcal{I}_t(\lambda))\right\}\\ &+ \text{meas}\left\{t \in [\tau(\eta_n K), \tau(\eta_n K + m)) : W(2\mathcal{I}_t(\lambda) - \eta_n K) > \sqrt{n}(1-\eta_n)K\right\}\\ &+ \text{meas}\{t \in [\tau(\eta_n K + m), \tau_- : W(\mathcal{I}_t(\lambda) + m) > \sqrt{n}(K + m - \mathcal{I}_t(\lambda))\}\\ &+ \text{meas}\left\{t \in [\tau_-, \tau_+] : W(\mathcal{I}_t(\lambda) + m) > \sqrt{n}(K + m - \mathcal{I}_t(\lambda))\right\}\\ &+ \text{meas}\left\{t \in [\tau_+, T] : W(\mathcal{I}_t(\lambda) + m) > \sqrt{n}(K + m - \mathcal{I}_t(\lambda))\right\}.\end{aligned} \tag{C.3.12}$$

It is best to focus on one term, i.e., one line of equation (C.3.12), at a time.

I. The first term in (C.3.12), namely,

$$\text{meas}\left\{t \in [0, \tau(\eta_n K)) : W(\mathcal{I}_t(\lambda)) > \sqrt{n}(K - \mathcal{I}_t(\lambda))\right\} \tag{C.3.13}$$

is evidently bounded from above by

$$\text{meas}\left\{t \in [0, \tau(\eta_n K)) : W(\mathcal{I}_t(\lambda)) > \sqrt{n}(1-\eta_n)K\right\},$$

and even more so by

$$\text{meas}\left\{t \in [0, T] : W(\mathcal{I}_t(\lambda)) > \sqrt{n}(1-\eta_n)K\right\}. \tag{C.3.14}$$

Let us take, using the nomenclature of Lemma C.2.5, $s_n = \sqrt{n}(1-\eta_n)K$, $\alpha_n \equiv 0$ and $f_n = \mathcal{I}(\lambda)$. By Assumption 3.3.4, the sequence $\{s_n\}$ is indeed divergent, and the other assumptions of Lemma C.2.5 are trivially satisfied. Hence, the expectation of the random variable in (C.3.14) disappears in the limit as $n \to \infty$. As the term in (C.3.14) is an upper bound on the term in (C.3.13), this term's expectation vanishes in the limit, as well.

II. The form of the second term allows us to use Lemma C.2.5 and Assumption 3.3.4 directly, concluding that its expectation too vanishes in the limit.

III. This time we apply Lemma C.2.6 to show that the third term's expectation converges to zero, as well.

IV. This is the first term whose limit is not entirely trivial. For all $t \in [\tau((K+m)-), \tau(K+m)]$, we have that $\mathcal{I}_t(\lambda) = K + m$, so that the fourth term in (C.3.12) can be rewritten as

$$\begin{aligned}&\text{meas}\{t \in [\tau((K+m)-), \tau(K+m)] : W(K+2m) > 0\}\\ &\quad = (\tau(K+m) - \tau((K+m)-))\mathbf{1}_{\{W(K+2m)>0\}}.\end{aligned}$$

Recognizing the result as a part of the random variable $\hat{J}^*$, we leave it in the present form.

V. By means of Corollary C.3.4, the final term converges in expectation to $(T - \tau(K+m))$.

Trivial algebra wraps up the proof. $\square$

C.4 Asymptotic Optimality Analysis Miscellany

C.4.1 Absence of Regulation

Lemma 3.4.4 establishes a lower bound on the cost aggregated on the segment $[\tau(\gamma_n K + m), T]$. Next, we wish to estimate, for all n, the probability of $\frac{1}{n}Q_t^{(n)}(\tilde{\mu}_n)$ ever exceeding the threshold K before time $\tau(\gamma_n K + m)$.

For each $n \in \mathbb{N}$, on the interval $[0, \tau(\gamma_n K)]$, the process $X^{(n)}(\tilde{\mu}_n)$ is non-negative, and so we have $Q_t^{(n)}(\tilde{\mu}_n) = X_t^{(n)}(\tilde{\mu}_n)$ on that segment. On the interval $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$, the potential departure process $N^-(n\mathcal{I}(\tilde{\mu}_n)) = N^-(n(\mathcal{I}(\lambda) - \gamma_n K))$ and the arrival process $N^+(n\mathcal{I}(\lambda))$ have the same instantaneous rates. Hence, conditionally on the behavior of the system on $[0, \tau(\gamma_n K)]$, the netput process $X^{(n)}(\tilde{\mu}_n)$ is centered at $N^+(n\gamma_n K)$, as its increments during the interval $[\tau(\gamma_n K), \tau(\gamma_n K+m)]$ have mean zero. Moreover, the deviation of the random variable $\frac{1}{n}Q^{(n)}{}_{\tau(\gamma_n K)}(\tilde{\mu}_n)$ from its mean $\gamma_n K$ will vanish in the limit, in the strong sense.

Lemma C.4.1. *As n tends to infinity,*

$$\left|\frac{1}{n}Q^{(n)}{}_{\tau(\gamma_n K)}(\tilde{\mu}_n) - \gamma_n K\right| \longrightarrow 0, \ a.s.$$

Proof. For any n and $t \leq \tau(\gamma_n K)$, by (3.1.1), we have

$$Q_t^{(n)}(\tilde{\mu}_n) = N^+(n\mathcal{I}_t(\lambda)).$$

In particular, using the definition of τ in (3.2.7), we get

$$Q^{(n)}{}_{\tau(\gamma_n K)}(\tilde{\mu}_n) = N^+(n\mathcal{I}_{\tau(\gamma_n K)}(\lambda)) = N^+(n\gamma_n K).$$

Therefore, for all n, we have

$$\left|\frac{1}{n}Q^{(n)}{}_{\tau(\gamma_n K)}(\tilde{\mu}_n) - \gamma_n K\right| = \left|\frac{1}{n}N^+(n\gamma_n K) - \gamma_n K\right|.$$

By the Strong Law of Large Numbers, the right-hand side tends to zero as $n \to \infty$, almost surely. $\square$

Lemma C.4.2. *For every $n \in \mathbb{N}$, the process $\{X_t^{(n)}(\tilde{\mu}_n); \tau(\gamma_n K) \leq t \leq \tau(\gamma_n K + m)\}$ is a martingale. Moreover, $\mathbb{E}[X_t^{(n)}(\tilde{\mu}_n)] = n\gamma_n K$ for every $t \in [\tau(\gamma_n K), \tau(\gamma_n K + m)]$.*

Proof. Integrability of all the processes involved is trivially satisfied.

Let us temporarily fix the index n, and consider $s \leq t$, both elements of $[\tau(\gamma_n K), \tau(\gamma_n K+m)]$. We have

$$\mathbb{E}[X_t^{(n)}(\tilde{\mu}_n) \,|\, X_s^{(n)}(\tilde{\mu}_n)] = \mathbb{E}[X_t^{(n)}(\tilde{\mu}_n) - X_s^{(n)}(\tilde{\mu}_n) \,|\, X_s^{(n)}(\tilde{\mu}_n)] + X_s^{(n)}(\tilde{\mu}_n).$$

Since N^+ and N^- are independent and, being Poisson processes, have independent increments the above yields

$$\begin{aligned}\mathbb{E}[X_t^{(n)}(\tilde{\mu}_n) \,|\, X_s^{(n)}(\tilde{\mu}_n)] = \mathbb{E}[&N^+(n\mathcal{I}_t(\lambda)) - N^+(n\mathcal{I}_s(\lambda))] \\ &- \mathbb{E}[N^-(n\mathcal{I}_t(\tilde{\mu}_n)) - N^-(n\mathcal{I}_s(\tilde{\mu}_n))] + X_s^{(n)}(\tilde{\mu}_n).\end{aligned}$$

Using the definition of $\tilde{\mu}_n$, we can rewrite the last equality as

$$\begin{aligned}\mathbb{E}[X_t^{(n)}(\tilde{\mu}_n) \,|\, X_s^{(n)}(\tilde{\mu}_n)] &= n\mathcal{I}_t(\lambda) - n\mathcal{I}_s(\lambda) - (n\mathcal{I}_t(\lambda) - n\mathcal{I}_s(\lambda)) + X_s^{(n)}(\tilde{\mu}_n) \\ &= X_s^{(n)}(\tilde{\mu}_n).\end{aligned}$$

Due to the martingality of the process $X^{(n)}(\tilde{\mu}_n)$, its expectation stays constant on the interval $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$ and equal to

$$\mathbb{E}[X_{\tau(\gamma_n K)}^{(n)}(\tilde{\mu}_n)] = \mathbb{E}[N^+(n\mathcal{I}_{\tau(\gamma_n K)}(\lambda))] = n\gamma_n K.$$

□

Next, we provide an upper bound for the probability of $\frac{1}{n}X^{(n)}(\tilde{\mu}_n)$ ever sinking to zero on the segment $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$.

Proposition C.4.3. *There exists a constant $C > 0$, such that for all $n \in \mathbb{N}$,*

$$\mathbb{P}\left[\inf_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) \leq 0\right] \leq \frac{C}{n^2\gamma_n^4}, \tag{C.4.1}$$

with the infimum taken across the segment $[0, \tau(\gamma_n K + m)]$.

Proof. From the definition of service disciplines $\tilde{\mu}_n$, we conclude that it suffices to consider the supremum in (C.4.1) merely over $[\tau(\gamma_n K), \tau(\gamma_n K \; + \; m)]$.

Let us temporarily fix the index n. For all $t \in [\tau(\gamma_n K), \tau(\gamma_n K + m)]$, $\mathcal{I}_t(\lambda) \leq \gamma_n K + m$ and $\mathcal{I}_t(\lambda) - \mathcal{I}_t(\tilde{\mu}_n) = \gamma_n K$. Hence, using Lemma C.2.8, we arrive at the inequality

$$\begin{aligned}\mathbb{P}\left[\inf_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) \leq 0\right] &= \mathbb{P}\left[\inf_t \frac{1}{n}[N^+(n\mathcal{I}_t(\lambda)) - N^-(n\mathcal{I}_t(\tilde{\mu}_n))] \leq 0\right] \\ &\leq \frac{8}{n^4\gamma_n^4K^4}[3n^2((\gamma_n K + m)^2 + m^2) + n(\gamma_n K + 2m)] \\ &\leq \frac{8}{n^2\gamma_n^4K^4}[3((K + m)^2 + m^2) + (K + 2m)].\end{aligned}$$

Taking $C = \frac{8}{K^4}[3((K+m)^2 + m^2) + (K + 2m)]$ completes the proof. □

The next corollary is the final claim establishing absence of reflection under Assumption 3.4.5(i).

Corollary C.4.4. *Let the sequence* $\{\gamma_n\}$ *be such that* $\sum_{n=1}^{\infty} \frac{1}{n^2\gamma_n^4} < \infty$. *Then*

$$\mathbb{P}\left[\inf_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) \leq 0, i.o.\right] = 0,$$

with the infimum taken over $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$. *In other words,* $\frac{1}{n}X_t^{(n)}(\tilde{\mu}_n) > 0$ *and, therefore,* $\frac{1}{n}X_t^{(n)}(\tilde{\mu}_n) = \frac{1}{n}Q_t^{(n)}(\tilde{\mu}_n)$*, for all* $t \in [0, T]$*, for all but finitely many* n*, almost surely.*

Proof. By the assumption, we have that

$$\sum_{n=1}^{\infty} \mathbb{P}\left[\inf_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) \leq 0\right] \leq C \sum_{n=1}^{\infty} \frac{1}{n^2\gamma_n^4} < \infty.$$

The corollary is then a straightforward consequence of the Borel-Cantelli Lemma. □

C.4.2 Absence of Penalty

Lemma C.4.5. *Let the sequence* $\{\gamma_n\}$ *satisfy Assumption 3.4.5, and take the constant* C *to be defined by*

$$C = \frac{8}{K^4}[3((K+m)^2 + m^2) + (K + 2m)].$$

Then, for all $n \in \mathbb{N}$*, the inequality*

$$\mathbb{P}\left[\sup_t \frac{1}{n} Q_t^{(n)}(\tilde{\mu}_n) > K\right] \leq \mathbb{P}\left[\sup_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) > K\right] + C\frac{1}{n^2\gamma_n^4} \tag{C.4.2}$$

holds, with the suprema taken over the segment $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$.

Proof. To simplify the exposition of this proof, we first introduce the following random variables

$$Q_*^{(n)} = \sup\{\frac{1}{n} Q_t^{(n)}(\tilde{\mu}_n) : \tau(\gamma_n K) \leq t \leq \tau(\gamma_n K + m))\},$$
$$X_*^{(n)} = \sup\{\frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) : 0 \leq t \leq \tau(\gamma_n K + m))\}, \text{ for every } n \in \mathbb{N}.$$

With this notation the desired inequality (C.4.2) is transformed into

$$\mathbb{P}\left[Q_*^{(n)} > K\right] \leq \mathbb{P}\left[X_*^{(n)} > K\right] + C\frac{1}{n^2\gamma_n^4}. \tag{C.4.3}$$

By the definition of the one-sided regulator map, in the event that $X_*^{(n)} > 0$, we necessarily have $X_*^{(n)} = Q_*^{(n)}$. Hence,

$$\mathbb{P}[X_*^{(n)} \neq Q_*^{(n)}] \leq \mathbb{P}[X_*^{(n)} \leq 0], \text{ for all } n.$$

Combined with the fact that $\{\gamma_n\}$ must satisfy Assumption 3.4.5, and the claim of Proposition C.4.3, the last display implies that

$$\mathbb{P}[X_*^{(n)} \neq Q_*^{(n)}] \leq C\frac{1}{n^2\gamma_n^4} \text{ for all } n. \tag{C.4.4}$$

Now we proceed to expand the left-hand side of (C.4.3) as

$$\begin{aligned}\mathbb{P}\left[Q_*^{(n)} > K\right] &= \mathbb{P}\left[Q_*^{(n)} > K,\, X_*^{(n)} = Q_*^{(n)}\right] + \mathbb{P}\left[Q_*^{(n)} > K,\, X_*^{(n)} \neq Q_*^{(n)}\right] \\ &\leq \mathbb{P}\left[X_*^{(n)} > K,\, X_*^{(n)} = Q_*^{(n)}\right] + \mathbb{P}\left[X_*^{(n)} \neq Q_*^{(n)}\right].\end{aligned}$$

Using (C.4.4), we get

$$\mathbb{P}\left[Q_*^{(n)} > K\right] \leq \mathbb{P}\left[X_*^{(n)} > K\right] + C\frac{1}{n^2\gamma_n^4}.$$

□

Proposition C.4.6. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5, and take the constant C to be defined as*

$$C = \frac{8}{K^4}[3((K+m)^2 + m^2) + (K+2m)].$$

Then, for all $n \in \mathbb{N}$, the inequality

$$\mathbb{P}\left[\sup_t \frac{1}{n}X_t^{(n)}(\tilde{\mu}_n) > K\right] \leq C\frac{1}{n^2(1-\gamma_n)^4} \tag{C.4.5}$$

holds, with the supremum taken over the segment $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$.

Proof. The proof mimics the steps of the proof of Proposition C.4.3. For all n, we can rewrite the left-hand side of (C.4.5) and rearrange the terms in the obtained expression to get

$$\begin{aligned}&\mathbb{P}\Big[\sup_t \frac{1}{n}X_t^{(n)}(\tilde{\mu}_n) > K\Big] \\ &\quad = \mathbb{P}\left[\sup_t \frac{1}{n}X_t^{(n)}(\tilde{\mu}_n) - \gamma_n K > K - \gamma_n K\right] \\ &\quad = \mathbb{P}\left[\sup_t \left(\frac{1}{n}[N^+(n\mathcal{I}_t(\lambda)) - N^-(n\mathcal{I}_t(\tilde{\mu}_n))] - \gamma_n K\right)^4 > (1-\gamma_n)^4 K^4\right].\end{aligned}$$

Combining Lemmas C.2.9 and C.4.2 with the above inequality, we obtain

$$\mathbb{P}\Big[\sup_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) > K\Big] \leq \frac{\mathbb{E}[(n\gamma_n K - X^{(n)}{}_{\tau(\gamma_n K+m)}(\tilde{\mu}_n))^4]}{n^4(1-\gamma_n)^4 K^4}. \tag{C.4.6}$$

Now, we do some simple algebra. The compensated netput process at time $\tau(\gamma_n K + m)$ can be expanded as

$$X^{(n)}{}_{\tau(\gamma_n K+m)}(\tilde{\mu}_n) - n\gamma_n K = (N^+(n\gamma_n K + nm) - n\gamma_n K - nm) - (N^-(nm) - nm).$$

The fourth centered moments of the number of arrivals and potential number of departures are evaluated using Lemma C.2.7 and then bounded from above. We obtain

$$\begin{aligned}\mathbb{E}[(N^+(n\gamma_n K + nm) - n\gamma_n K - nm)^4] &= 3n^2(\gamma_n K + m)^2 + n(\gamma_n K + m)\\ &\leq 3n^2[(K+m)^2 + (K+m)],\end{aligned} \tag{C.4.7}$$

and

$$\mathbb{E}[(N^-(nm) - nm)^4] = 3n^2 m^2 + nm \leq 3n^2[m^2 + m]. \tag{C.4.8}$$

Now we use Lemma D.5.1 and inequalities in (C.4.7) and (C.4.8) to obtain the following upper bound for the expression in (C.4.6):

$$\mathbb{P}\left[\sup_t \frac{1}{n} X_t^{(n)}(\tilde{\mu}_n) > K\right] \leq C\frac{1}{n^2(1-\gamma_n)^4}. \tag{C.4.9}$$

This is exactly the claim we are proving. □

We combine the last two results to get the following proposition.

Proposition C.4.7. *Let the sequence $\{\gamma_n\}$ satisfy Assumption 3.4.5, and let the constant C be defined by*

$$C = \frac{8}{K^4}[3((K+m)^2 + m^2) + (K+2m)].$$

Then, for all $n \in \mathbb{N}$, the inequality

$$\mathbb{P}\left[\sup_t \frac{1}{n} Q_t^{(n)}(\tilde{\mu}_n) > K\right] \leq C\left[\frac{1}{n^2(1-\gamma_n)^4} + \frac{1}{n^2\gamma_n^4}\right] \tag{C.4.10}$$

holds, with the supremum taken over the segment $[\tau(\gamma_n K), \tau(\gamma_n K + m)]$.

Proof. The claim is a straightforward consequence of Lemma C.4.5 and Proposition C.4.6. □

C.4.3 A Special case

Lemma C.4.8. *Let us assume that there exists a neighborhood I of $\tau(K+m)$ such that $\lambda_t > 0$ for all $t \in I$. Then*

$$\tau_*^{(n)} \to \tau(K+m), \text{ a.s., as } n \to \infty.$$

Proof. According to Theorem 9.6.1 in [Whi02a], we have that

$$\frac{1}{n} N^+(n\mathcal{I}(\lambda)) \to \mathcal{I}(\lambda), a.s., \tag{C.4.11}$$

uniformly on compacts. Let us recall the definition (3.4.3) of the processes $\tau^{(n)}$. They are given as (right-continuous) inverses of the normalized arrival processes that appear on the left-hand side in (C.4.11). Using Theorem 13.6.3 in [Whi02b], we conclude that necessarily $\tau^{(n)} \to \tau$, almost surely in M_1-topology.

In fact, more is true on the interval I under the assumption of positivity of λ. In this case, the function $\mathcal{I}(\lambda)$ is strictly increasing on I, and so τ becomes the proper inverse of $\mathcal{I}(\lambda)$ and is continuous on the restricted domain I (as opposed to being merely a right-continuous inverse). This holds by Lemma 13.6.5 in [Whi02b]. This fact implies that the almost sure convergence of the sequence of random processes $\tau^{(n)}$ defined in (3.4.3) towards τ is uniform on this domain (see p.82 in [Whi02b]).

By the Strong Law of Large Numbers, we have

$$\frac{1}{n} N^-(nm) \to m, \text{ a.s.} \tag{C.4.12}$$

So for all n, we see that

$$\begin{aligned} |\tau_*^{(n)} - \tau(K+m)| &= \left|\tau^{(n)}\left(K + \frac{1}{n}N^-(nm)\right) - \tau(K+m)\right| \\ &\le \left|\tau^{(n)}\left(K + \frac{1}{n}N^-(nm)\right) - \tau\left(K + \frac{1}{n}N^-(nm)\right)\right| \\ &\quad + \left|\tau\left(K + \frac{1}{n}N^-(nm)\right) - \tau(K+m)\right|. \end{aligned} \tag{C.4.13}$$

For large enough n, we have that $K + \frac{1}{n}N^-(nm) \in I$. Hence, the first term on the right-hand side of (C.4.13) vanishes in the limit since $\tau^{(n)} \to \tau$ uniformly on I, and the argument $K + \frac{1}{n}N^-(nm)$ is eventually in I. The second term on the right-hand side of (C.4.13) disappears in the limit due to the continuity of τ on I and the convergence in (C.4.12). □

When combined with Theorem 3.4.8, the last lemma yields the following result.

Corollary C.4.9. *Assume that there exists a neighborhood I of $\tau(K+m)$ such that $\lambda_t > 0$ for all $t \in I$. Then we have*

$$J^{(n)}(\tilde{\mu}_n) \to J^*, \text{ a.s., as } n \to \infty,$$

where J^* is the optimal value of the performance measure $\bar{J}$ from the fluid limit analysis, and is defined in (3.2.8).

C.5 Comparison of Classes

The first subsection is a necessary disclaimer regarding the nature of comparison between asymptotically optimal and second order optimal sequences.

C.5.1 Asymptotically Optimal vs. Second Order Optimal

It is not surprising that the class of sequences we obtained as second order optimal in the sense of Definition 3.2.2 is similar to the class of disciplines that we obtained as asymptotically optimal in the Section 3.4. Comparing Assumptions 3.3.4 and 3.4.5, it is obvious that the conditions of Assumption 3.3.4 are more relaxed that the ones of Assumption 3.4.5. In other words, the class of asymptotically optimal sequences is a subset of the class of second order optimal sequences of service disciplines. Before we explore their relationship further, we should be careful and state that unfortunately in both cases the conditions are merely *sufficient* for optimality. It would make more sense to be comparing the classes of optimal sequences of service disciplines were they completely described, i.e., had a criterion both necessary and sufficient been provided. However, in view of the natural construction of both classes, let us continue.

Now that the most obvious drawback has been pointed out, the next section is dedicated to the discrepancy between Definitions 3.3.1 and 3.4.2.

C.5.2 "Almost Surely" vs. "in Expectation"

It would not be difficult to argue that an almost sure result in a control problem is preferable to a result given in expectation. If we accept the sequence of uniformly accelerated systems as a reasonable tool for grasping the actual system, an almost surely asymptotically optimal discipline simply means that no matter what the state of the world, our control is optimal in a pathwise sense. A result in expectation, on the other hand, includes a possibility of a certain state of the world working extremely against us. In other words, our performance may be suboptimal.

Let us recall that the asymptotic expansion in (2.2.4) holds merely in the distributional sense. Hence, the requirement in Definition 3.3.1 for second order optimal sequences to outperform any other deterministic admissible sequence *in expectation* is indeed a sensible one. Also, due to the nature of the paths of Brownian motion (they do not have monotone stretches), all chances of stochastic comparison between paths of the approximating processes $\bar{Q}(\mu) + \frac{1}{\sqrt{n}}\hat{Q}(\mu)$, as the service discipline μ is varied, is lost. On the other hand, the non-homogenous Poisson processes in the prelimit systems $Q^{(n)}(\mu)$ could be compared for different service disciplines μ. This allowed us to require asymptotic optimality in the almost sure sense, as stated in Definition 3.4.2.

In order to make the comparison of optimal sequences in the asymptotic and the second order sense more reasonable, let us introduce the following, more relaxed, notion of asymptotic optimality.

Definition C.5.1. An admissible sequence $\{\mu_n^*\}$ will be called *weakly asymptotically optimal* for the sequence of performance measures $\{J^{(n)}\}$ given in (4.3.2), if

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^{(n)}(\mu_n^*)] \geq 0,$$

for any other admissible sequence $\{\mu_n\}$.

Obviously, an analogue of Lemma 3.3.2 holds, so let us determine a suitable constant J^*, such that

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^*] \geq 0,$$

for any sequence $\{\mu_n\}$ of admissible service disciplines. Let $J^* = \hat{J}^*$, as introduced in (3.3.6). Then the following lemma holds true.

Lemma C.5.2. *For all admissible sequences* $\{\mu_n\}$,

$$\liminf_{n\to\infty} \mathbb{E}[J^{(n)}(\mu_n) - J^*] \geq 0.$$

Proof. For any n, and any choice of μ_n, by Fubini's theorem,

$$\mathbb{E}[J^{(n)}(\mu_n)] = \int_0^T \mathbb{P}\left[\frac{1}{n}Q_t^{(n)}(\mu_n) > K\right] dt \geq \int_{\tau((K+m)-)}^T \mathbb{P}\left[\frac{1}{n}Q_t^{(n)}(\mu_n) > K\right] dt. \tag{C.5.1}$$

For all t, using the monotonicity of Poisson processes, we have

$$\begin{aligned} \frac{1}{n}Q_t^{(n)}(\mu_n) \geq \frac{1}{n}X_t^{(n)}(\mu_n) &= \frac{1}{n}[N^+(n\mathcal{I}_t(\lambda)) - N^-(n\mathcal{I}_t(\mu_n))] \\ &\geq \frac{1}{n}[N^+(n\mathcal{I}_t(\lambda)) - N^-(nm)]. \end{aligned} \tag{C.5.2}$$

Let us fix a time t, such that $\mathcal{I}_t(\lambda) = K + m$. Then

$$\begin{aligned} \mathbb{P}\left[\frac{1}{n}Q_t^{(n)}(\mu_n) > K\right] &\geq \mathbb{P}\left[\frac{1}{n}[N^+(n(K+m)) - N^-(nm)] > K\right] \\ &= \mathbb{P}\left[\sqrt{n}\left(\frac{1}{n}N^+(n(K+m)) - (K+m)\right) - \sqrt{n}\left(\frac{1}{n}N^-(nm) - m\right)\right] > 0]. \end{aligned}$$

By the Central Limit Theorem and independence of N^+ and N^-, this probability converges to $\frac{1}{2}$.

On the other hand, for all $t > \tau(K+m)$, by the Strong Law of Large Numbers, the right-hand side of (C.5.2) converges almost surely towards $\mathcal{I}_t(\lambda) - m > K$. Therefore, as $n \to \infty$, for any fixed $t > \tau(K+m)$,

$$\mathbb{P}\left[\frac{1}{n} Q_t^{(n)}(\mu_n) > K\right] \to 1.$$

Performing the integration indicated in (C.5.1) completes the proof. □

Next, let us verify that the class of admissible sequences of service disciplines defined by

$$\hat{\mu}_n = \lambda \mathbf{1}_{[\tau(\eta_n K), \tau(\eta_n K + m)]}, \tag{C.5.3}$$

for sequences $\{\eta_n\}$ satisfying

$$n(1-\eta_n)^2 \longrightarrow \infty, \text{ as } n \to \infty, \tag{C.5.4}$$

are not only second-order optimal, but also weakly asymptotically optimal in the sense of Definition C.5.1, provided that we additionally assume

$$\sum_{n=1}^{\infty} \frac{1}{n^2 \eta_n^4} < \infty. \tag{C.5.5}$$

The assumption (C.5.5) is in place to ensure (asymptotical) absence of reflection, and assumption (C.5.4) allows the system to (asymptotically) avoid unnecessary transitions over the threshold. We have already encountered assumption (C.5.5), with somewhat different notation, as part of Assumption 3.4.5.

Lemma C.5.3. *Let $\{\eta_n\}$ satisfy assumptions (C.5.4) and (C.5.5). Define the admissible sequence $\{\hat{\mu}_n\}$ as in (3.3.7). Then, the following claims hold true:*

(i) for all t such that $\mathcal{I}_t(\lambda) = K + m$,

$$\mathbb{P}\Big[\frac{1}{n} Q_t^{(n)}(\hat{\mu}_n) > K\Big] \to \frac{1}{2}, \text{ as } n \to \infty;$$

(ii) for all t such that $\mathcal{I}_t(\lambda) < K + m$,

$$\mathbb{P}\left[\frac{1}{n} Q_t^{(n)}(\hat{\mu}_n) > K\right] \to 0, \text{ as } n \to \infty. \tag{C.5.6}$$

Proof. We prove the proposed statements one at a time.

Statement (i). Let us temporarily fix an instant t such that $\mathcal{I}_t(\lambda) = K + m$. For all n, from the definition of $\hat{\mu}_n$, we conclude that $\mathcal{I}_t(\hat{\mu}_n) = m$. Therefore, the claim

$$\frac{1}{n} X_t^{(n)}(\hat{\mu}_n) > K \tag{C.5.7}$$

is equivalent to

$$\sqrt{n}\left(\frac{1}{n}N^{+}(n(K+m)) - (K+m)\right) - \sqrt{n}\left(\frac{1}{n}N^{-}(nm) - m\right) > 0. \tag{C.5.8}$$

By the Central Limit Theorem, the random variable on the left-hand side of (C.5.8) converges to a normally distributed random variable centered at zero. Hence, we get

$$\mathbb{P}\left[\sqrt{n}\left(\frac{1}{n}N^{+}(n(K+m)) - (K+m)\right) > \sqrt{n}\left(\frac{1}{n}N^{-}(nm) - m\right)\right] \to \frac{1}{2}.$$

Using the equivalence of (C.5.7) and (C.5.8), we conclude that

$$\lim_{n\to\infty} \mathbb{P}\left[\frac{1}{n} X_t^{(n)}(\hat{\mu}_n) > K\right] = \frac{1}{2}. \tag{C.5.9}$$

Assumption (C.5.5) and Corollary C.4.4 imply that

$$\mathbb{P}[X_t^{(n)}(\hat{\mu}_n) = Q_t^{(n)}(\hat{\mu}_n),\, ev.] = 1. \tag{C.5.10}$$

Using Lemma C.2.2 along with conditions (C.5.9) and (C.5.10), we verify the claim of the lemma.

Statement (ii). Let $\{\eta_{n_k}\}_k$ be an arbitrary subsequence of the sequence $\{\eta_n\}$. Since all the elements of $\{\eta_{n_k}\}_k$ are bounded between 0 and 1, this subsequence has a further subsequence $\{\eta_{n_{k_l}}\}_l$ which is convergent. For typographical convenience, let us denote this subsequence by $\{\nu_l\}$ and call its limit ν. Moreover, we define the sequence of natural numbers $\{c_l\}$ by $c_l = n_{k_l}$, for all indices l. Assumption (C.5.4) implies that

$$c_l(1-\nu_l)^2 \to \infty, \text{ as } l \to \infty. \tag{C.5.11}$$

The sequence $\{\nu_l\}$ through the definition in (3.3.7) generates a sequence of service disciplines $\{\hat{\mu}_{c_l}\}_l$. The corresponding sequence of cumulative services in the system is then given by

$$\mathcal{I}(\hat{\mu}_{c_l}) = (\mathcal{I}(\lambda) - \nu_l K)^{+} \wedge m. \tag{C.5.12}$$

The sequence of functions in (C.5.12) has as its uniform limit $(\mathcal{I}(\lambda) - \nu K)^{+} \wedge m$.

Now, for any fixed instant t satisfying $\mathcal{I}_t(\lambda) < K + m$, the Strong Law of Large Numbers implies that

$$\frac{1}{c_l} X_t^{(c_l)}(\hat{\mu}_{c_l}) \to \bar{x}_t := \mathcal{I}_t(\lambda) - (\mathcal{I}_t(\lambda) - \nu K)^+ \wedge m. \tag{C.5.13}$$

The value on the right-hand side of (C.5.13) can be written more conveniently as

$$\bar{x}_t = \begin{cases} \mathcal{I}_t(\lambda) & \text{for } t < \tau(\eta K) \\ \eta K & \text{for } \tau(\eta K) \leq t < \tau(\eta K + m) \\ \mathcal{I}_t(\lambda) - m & \text{for } \tau(\eta K + m) \leq t < \tau((K+m)-). \end{cases} \tag{C.5.14}$$

There are two cases that need to be considered separately here, depending on the value of ν.

Case 1. Let $\nu < 1$. Then for all $t < \tau((K+m)-)$, i.e., for all t such that $\mathcal{I}_t(\lambda) < K + m$, we have $\bar{x}_t < K$. Using the Portmanteau Theorem, we conclude that

$$\mathbb{P}\left[\frac{1}{c_l} X_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to \mathbb{P}[\bar{x}_t > K] = 0, \text{ as } l \to \infty. \tag{C.5.15}$$

Using Lemma C.2.2 and claims (C.5.10) and (C.5.15), we get

$$\mathbb{P}\left[\frac{1}{c_l} Q_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to 0, \text{ as } l \to \infty.$$

Case 2. In this case $\nu = 1$. So, $\nu_l \uparrow \nu$, as $l \to \infty$. Recalling the definition (3.2.7) of the mapping τ, we conclude that $\tau(\nu_l) \to \tau(K-)$, as $l \to \infty$.

For any $t < \tau(K-)$, the description of $\bar{x}_t$ given in (C.5.14) yields that $\bar{x}_t < K$. Hence, the Portmanteau theorem is applicable to get

$$\mathbb{P}\left[\frac{1}{c_l} X_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to \mathbb{P}[\bar{x}_t > K] = 0, \text{ as } l \to \infty. \tag{C.5.16}$$

In the exact same fashion as in the first case we conclude that

$$\mathbb{P}\left[\frac{1}{c_l} Q_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to 0, \text{ as } l \to \infty.$$

To the contrary, for all $t \in [\tau(K-), \tau((K+m)-)]$ we have $\bar{x}_t = K$, which renders the Portmanteau theorem inapplicable if we wish to reach a conclusion similar to the one in (C.5.16). However, we can proceed by noting that, for any l, simple algebra gives us

$$\begin{aligned} &\mathbb{P}\left[\frac{1}{c_l} X_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \\ &\quad = \mathbb{P}\left[\sqrt{c_l}\left(\frac{1}{c_l} X_t^{(c_l)}(\hat{\mu}_{c_l}) - (\mathcal{I}_t(\lambda) - \mathcal{I}_t(\hat{\mu}_{c_l}))\right) > \sqrt{c_l}(K - (\mathcal{I}_t(\lambda) - \mathcal{I}_t(\hat{\mu}_{c_l})))\right]. \end{aligned}$$

By the equality in (C.5.12), the last expression yields

$$\mathbb{P}\left[\frac{1}{c_l}X_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] = \mathbb{P}\left[\sqrt{c_l}\left(\frac{1}{c_l}X_t^{(c_l)}(\hat{\mu}_{c_l}) - (\mathcal{I}_t(\lambda) - \mathcal{I}_t(\hat{\mu}_{c_l}))\right) > \sqrt{c_l}(1-\nu_l)K\right]. \quad \text{(C.5.17)}$$

The Central Limit Theorem implies that

$$\sqrt{c_l}\left(\frac{1}{c_l}X_t^{(c_l)}(\hat{\mu}_{c_l}) - (\mathcal{I}_t(\lambda) - \mathcal{I}_t(\hat{\mu}_{c_l}))\right) \Rightarrow N(0, \Sigma_t), \text{ as } l \to \infty, \quad \text{(C.5.18)}$$

where Σ_t denotes the variance of the limiting normal distribution whose exact value is irrelevant for our purposes. Of course, this random variable almost surely attains finite values. Therefore, looking at (C.5.11), we realize that (C.5.17) and (C.5.18) together imply

$$\mathbb{P}\left[\frac{1}{c_l}X_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to \mathbb{P}[\bar{x}_t > K] = 0, \text{ as } l \to \infty.$$

So that, again, we have

$$\mathbb{P}\left[\frac{1}{c_l}Q_t^{(c_l)}(\hat{\mu}_{c_l}) > K\right] \to 0, \text{ as } l \to \infty.$$

We have just proven that for any subsequence of the sequence of numbers $\mathbb{P}[\frac{1}{n}Q_t^{(n)}(\hat{\mu}_n) > K]$ there exists a further *convergent* subsequence. Moreover, all of these subsubsequences converge to the same limit, namely, zero. It is a well known result that then the sequence $\mathbb{P}[\frac{1}{n}Q_t^{(n)}(\hat{\mu}_n) > K]$ itself must converge, and to the same limit. This completes the proof of claim (C.5.6). $\square$

Appendix D

The Tandem System - Auxiliary Results

D.1 Transition Probabilities

For the purposes of this section we extend the existing notation by introducing the functional $\mathcal{I}_\infty$ given as $\mathcal{I}_\infty(f) = \int_0^\infty f(s)\,ds$ for every integrable function $f : [0,\infty) \to \mathbb{R}$.

Lemma D.1.1. *Suppose that Y_1 and Y_2 are two independent unit Poisson processes on a common probability space. Let $\rho_1 : [0,\infty) \to (0,\infty)$ and $\rho_2 : [0,\infty) \to [0,\infty)$ be integrable functions such that $\mathcal{I}_\infty(\rho_2) > 0$.*

We define the stopping times σ_1 and σ_2 denoting the first jumps in the two time-changed Poisson processes as

$$\sigma_1 = \inf\{t \in [0,\infty) : Y_1(\mathcal{I}_t(\rho_1)) > 0\},$$
$$\sigma_2 = \inf\{t \in [0,\infty) : Y_2(\mathcal{I}_t(\rho_2)) > 0\}.$$

Then we have

$$\mathbb{P}[\sigma_1 < \sigma_2] = \int_0^\infty \left(1 - e^{-\mathcal{I}_s(\rho_1)}\right) \rho_2(s) e^{-\mathcal{I}_s(\rho_2)} ds.$$

Proof. In order to compute the probability of interest, we begin by identifying the distribution function of the stopping times σ_1 and σ_2, denoted by F_{σ_1} and F_{σ_2}, respectively. The stopping times are independent due to the independence of Y_1 and Y_2, so their distribution functions will uniquely determine the joint distribution function of the pair (σ_1, σ_2).

For any positive real number s, we have

$$1 - F_{\sigma_1}(s) = \mathbb{P}[\sigma_1 > s] = \mathbb{P}[Y_1(\mathcal{I}_s(\rho_1)) = 0] = e^{-\mathcal{I}_s(\rho_1)}.$$

Analogously, $F_{\sigma_2}(s) = 1 - e^{-\mathcal{I}_s(\rho_2)}$.

The rest is a simple computation yielding

$$\mathbb{P}[\sigma_1 < \sigma_2] = \int_0^\infty (1 - e^{-\mathcal{I}_s(\rho_1)})\rho_2(s)e^{-\mathcal{I}_s(\rho_2)}ds.$$

□

Remark D.1.1. One should note that the essential ingredient for the above proof is the fact that the given rates are not identically zero which prevents the trivial case of stopping times concentrating at the "graveyard" state of ∞.

The following corollary gives us the limiting result upon acceleration of the system from the previous lemma. Note that the rate of acceleration of ρ_1 is much larger than the one for ρ_2.

Corollary D.1.2. *Suppose that $\rho_2 : \mathbb{R}_+ \to \mathbb{R}_+$ is an integrable function such that $\mathcal{I}_\infty(\rho_2) > 0$ and define a sequence of rate functions $\rho_1^{(n)} \equiv \rho_2 g(n)/n$, where $g : \mathbb{N} \to (0, \infty)$ is such that $\frac{g(n)}{n} \to \infty$ as $n \to \infty$. We set the following two sequences of stopping times*

$$\sigma_1^{(n)} = \inf\{t \in [0, \infty) : Y_1(n\mathcal{I}_t(\rho_1^{(n)})) > 0\},$$
$$\sigma_2^{(n)} = \inf\{t \in [0, \infty) : Y_2(n\mathcal{I}_t(\rho_2)) > 0\}.$$

Then for every n, we have

$$\mathbb{P}[\sigma_1^{(n)} < \sigma_2^{(n)}] = 1 - e^{-n\mathcal{I}_\infty(\rho_2)} - \frac{1}{1 + \frac{g(n)}{n}}(1 - e^{-(n+g(n))\mathcal{I}_\infty(\rho_2)}),$$

where $\mathcal{I}_\infty(\rho_2) = \int_0^\infty \rho_2(s)\, ds$. Moreover, as $n \to \infty$

$$\mathbb{P}[\sigma_1^{(n)} < \sigma_2^{(n)}] \longrightarrow 1.$$

Proof. According to the previous lemma, for all n

$$\begin{aligned}\mathbb{P}[\sigma_1^{(n)} < \sigma_2^{(n)}] &= n \int_0^\infty (1 - e^{-n\mathcal{I}_s(\rho_1^{(n)})})\rho_2(s)e^{-n\mathcal{I}_s(\rho_2)}\, ds \\ &= n \int_0^\infty (1 - e^{-g(n)\mathcal{I}_s(\rho_2)})\rho_2(s)e^{-n\mathcal{I}_s(\rho_2)}\, ds \\ &= n \left[\int_0^\infty \rho_2(s)e^{-n\mathcal{I}_s(\rho_2)}\, ds - \int_0^\infty \rho_2(s)e^{-(n+g(n))\mathcal{I}_s(\rho_2)}\, ds\right].\end{aligned}$$

Evaluating the integrals in the last display, we obtain

$$\begin{aligned}\mathbb{P}[\sigma_1^{(n)} < \sigma_2^{(n)}] &= n \left[\frac{1}{n}(1 - e^{-n\mathcal{I}_\infty(\rho_2)}) - \frac{1}{n + g(n)}(1 - e^{-(n+g(n))\mathcal{I}_\infty(\rho_2)})\right] \\ &= 1 - e^{-n\mathcal{I}_\infty(\rho_2)} - \frac{1}{1 + \frac{g(n)}{n}}(1 - e^{-(n+g(n))\mathcal{I}_\infty(\rho_2)}).\end{aligned}$$

The right-hand side converges to 1, by the given assumption on the growth of g. □

Let us introduce a sequence $\{\kappa_n\}$ of positive integers. The next question we wish to answer is the one concerning the limiting behavior of the probabilities established above, but this time with an independent κ_n-tuple of systems from the previous corollary.

Corollary D.1.3. *Let us assume that for every n, we have a family of κ_n pairs of independent driving unit Poisson processes which are denoted by $Y_1^{(i)}$ and $Y_2^{(i)}$ for $i \leq \kappa_n$ and accommodated on a common probability space. Moreover, for every n and i, we define the random variables*

$$\sigma_1^{(n,i)} = \inf\{t \in [0,\infty) : Y_1^{(i)}(n\mathcal{I}_t(\rho_1^{(n,i)})) > 0\},$$
$$\sigma_2^{(n,i)} = \inf\{t \in [0,\infty) : Y_2^{(i)}(n\mathcal{I}_t(\rho_2^{(n,i)})) > 0\},$$

where $\rho_1^{(n,i)} = \rho_2^{(n,i)} g(n)/n$ and $\rho_2^{(n,i)} \in \mathbb{L}_+^1[0,T]$ are such that $\mathcal{I}_\infty(\rho_2^{(n,i)}) > \epsilon$ for every pair of indices (n,i) and some positive constant ϵ. We further assume that

$$\kappa_n \ln\left(\frac{g(n)}{g(n)+n}\right) \to 0 \text{ as } n \to \infty \tag{D.1.1}$$

and

$$\kappa_n e^{-n\epsilon} \to 0 \text{ as } n \to \infty. \tag{D.1.2}$$

Then, as $n \to \infty$,

$$\mathbb{P}[\forall i \leq \kappa_n : \sigma_1^{(n,i)} < \sigma_2^{(n,i)}] \longrightarrow 1.$$

Proof. Since for all n, the κ_n systems are assumed to be driven by independent random processes, we have

$$\mathbb{P}[\forall i \leq \kappa_n : \sigma_1^{(n,i)} < \sigma_2^{(n,i)}] = \prod_{i \leq \kappa_n} \mathbb{P}[\sigma_1^{(n,i)} < \sigma_2^{(n,i)}].$$

Due to the continuity of the logarithmic function, the claim of the corollary is now equivalent to

$$\sum_{i=1}^{\kappa_n} \ln(\mathbb{P}[\sigma_1^{(n,i)} < \sigma_2^{(n,i)}]) \longrightarrow 0, \text{ as } n \to \infty.$$

For every n, using Corollary D.1.2, we get the following expansion of the sum in the above display

$$\begin{aligned}\sum_{i=1}^{\kappa_n} \ln(\mathbb{P}[\sigma_1^{(n,i)} < \sigma_2^{(n,i)}]) &= \sum_{i\le\kappa_n} \ln\left(1 - e^{-n\mathcal{I}_\infty(\rho_2^{(n,i)})} - \frac{n}{n+g(n)}\left(1 - e^{-(n+g(n))\mathcal{I}_\infty(\rho_2^{(n,i)})}\right)\right) \\ &= \sum_{i\le\kappa_n} \ln\left(\frac{g(n)}{n+g(n)} - e^{-n\mathcal{I}_\infty(\rho_2^{(n,i)})}\left(1 - e^{-g(n)\mathcal{I}_\infty(\rho_2^{(n,i)})}\right)\right) \\ &\ge \sum_{i\le\kappa_n} \ln\left(\frac{g(n)}{n+g(n)} - e^{-n\mathcal{I}_\infty(\rho_2^{(n,i)})}\right) \\ &\ge \sum_{i\le\kappa_n} \ln\left(\frac{g(n)}{n+g(n)} - e^{-n\epsilon}\right) \\ &= \kappa_n \ln\left(\frac{g(n)}{n+g(n)} - e^{-n\epsilon}\right).\end{aligned} \tag{D.1.3}$$

Due to the concavity of the logarithmic function, the following inequality holds true

$$\kappa_n \ln\left(\frac{g(n)}{n+g(n)} - e^{-n\epsilon}\right) \ge \kappa_n \ln\left(\frac{g(n)}{n+g(n)}\right) - \kappa_n \frac{e^{-n\epsilon}}{\frac{g(n)}{n+g(n)} - e^{-n\epsilon}}.$$

Straight from assumptions (D.1.1) and (D.1.2), the left-hand side above vanishes as $n \to \infty$. □

Example D.1.4. The conditions on the growth of the functions κ_n and $g(n)$ enforced in the previous corollary require to be explained further, as well as given in a more tractable form.

Let us assume that a divergent sequence $\{\kappa_n\}$ is given by $\kappa_n = \lfloor n^{1+\epsilon} \rfloor$, for all n, so that $\kappa_n e^{-n\epsilon} \to 0$ is automatically satisfied for any $\epsilon > 0$. Then, define g so that it satisfies $\frac{g(n)}{n} \to \infty$ as $n \to \infty$ in the following way:

$$g(n) = n \frac{e^{-\frac{1}{n^{1+2\epsilon}}}}{1 - e^{-\frac{1}{\kappa_n n^{1+2\epsilon}}}},$$

for the constant $\epsilon > 0$ from the corollary. It can be easily verified that $\kappa_n \ln(\frac{g(n)}{g(n)+n}) \to 0$.

D.2 Some Useful Integrals

Let f denote an integrable function on $[0, T]$, and let its integral be denoted by $F = \mathcal{I}(f)$. We assume that F is a nonnegative function. Furthermore, let c be a positive constant. We want to evaluate the following three integrals for every $t \in [0, T]$:

1. $\int_0^t f_s \mathbf{1}_{\{F_s < c\}}\, ds$;
2. $\int_0^t f_s \mathbf{1}_{\{F_s > c\}}\, ds$;

3. $\int_0^t f_s \mathbf{1}_{\{F_s=c\}}\, ds$.

Let us dedicate our attention to one integral at the time.

Integral 1. Since F is given as an integral, it is necessarily a continuous function. Therefore, the set $A = \{s \in [0,T] : F_s < c\}$ is an open subset of $[0,T]$. Moreover, since $[0,T]$ can be viewed as a subspace of $\mathbb{R}$, it is necessary that there exists an open set O in $\mathbb{R}$ such that $A = O \cap [0,T]$. Every open set on the real line can be expressed as a countable union of disjoint open intervals. In particular, this means that there exists a sequence of mutually disjoint intervals $\{(\alpha_i, \beta_i)\}$ such that $O = \cup_{i\in\mathbb{N}}(\alpha_i, \beta_i)$. Since $F_0 = 0$, we have that $0 \in A$. Therefore, without loss of generality we can assume that α_1 and β_1 are such that $\alpha_1 < 0 < \beta_1$. Also, the set A can now be displayed as $A = ([0,\beta_1) \cup \cup_{i\geq 2}(\alpha_i, \beta_i)) \cap [0,T]$.

Another consequence of the continuity of F is that for all $i > 1$ such that $\beta_i \leq T$, we have $F_{\alpha_i} = F_{\beta_i} = c$, and if $\alpha_i \leq T$ and $\beta_i > T$, then $F_{\alpha_i} = c$ (there can be at most one of the latter i's, since the intervals are assumed to be disjoint).

Now, the integral in question can be rewritten as follows:

$$\begin{aligned}
\int_0^t f_s \mathbf{1}_{\{F_s<c\}}\, ds &= \int_0^t f_s \mathbf{1}_A\, ds \\
&= \int_0^{\beta_1} f_s\, ds + \int_0^t f_s \sum_{i\geq 2} \mathbf{1}_{(\alpha_i,\beta_i)}(s)\, ds \\
&= c + \int_0^t f_s \lim_{N\to\infty} \sum_{i=2}^{N} \mathbf{1}_{(\alpha_i,\beta_i)}(s)\, ds.
\end{aligned} \tag{D.2.1}$$

So, by the Lebesgue's Dominated Convergence Theorem, we obtain

$$\int_0^t f_s \mathbf{1}_{\{F_s<c\}}\, ds = c + \lim_{N\to\infty} \int_0^t f_s \sum_{i=2}^{N} \mathbf{1}_{(\alpha_i,\beta_i)}(s)\, ds. \tag{D.2.2}$$

For every index N, Fubini's theorem (or its simple form - the additivity of the integral) allows us to transform the integral on the right-hand side of the last display by interchanging the order of summation and integration. We obtain the following.

$$\int_0^t f_s \sum_{i=2}^{N} \mathbf{1}_{(\alpha_i,\beta_i)}(s)\, ds = \sum_{i=2}^{N} \int_0^t f_s \mathbf{1}_{(\alpha_i,\beta_i)}(s)\, ds = \sum_{i=2}^{N} \int_{\alpha_i\wedge t}^{\beta_i\wedge t} f_s\, ds = \sum_{i=2}^{N} (F_{\beta_i\wedge t} - F_{\alpha_i\wedge t}).$$

For all $i \in \mathbb{N}$ such that $\beta_i \leq t$, the term in the above summation vanishes. The same happens for all i such that $\alpha_i \geq t$. By the disjointness assumption imposed on the sequence of open intervals at hand, there can be at most one index, say i^*, such that $t \in (\alpha_{i^*}, \beta_{i^*})$, and this is only the case if $F_t < c$.

Combining this conclusion with the expression (D.2.2), we obtain

$$\int_0^t f_s \mathbf{1}_{\{F_s<c\}}\, ds = \begin{cases} c + (F_t - c) & \text{if } F_t < c \\ c & \text{if } F_t \geq c \end{cases} \tag{D.2.3}$$
$$= F_t \wedge c.$$

Integral 2. One can quite effortlessly retrace the steps in the calculation for the previous case. The result is the following:

$$\int_0^t f_s \mathbf{1}_{\{F_s>c\}}\, ds = (F_t - c)^+. \tag{D.2.4}$$

Integral 3. Here we will simply combine the results of the above two calculations to obtain:

$$\begin{aligned} \int_0^t f_s \mathbf{1}_{\{F_s=c\}}\, ds &= \int_0^t f_s\, ds - \int_0^t f_s \mathbf{1}_{\{F_s<c\}}\, ds - \int_0^t f_s \mathbf{1}_{\{F_s=c\}}\, ds \\ &= F_t - F_t \wedge c - (F_t - c)^+ \\ &= \begin{cases} F_t - F_t & \text{if } F_t < c \\ F_t - c - F_t + c & \text{if } F_t \geq c \end{cases} \\ &= 0. \end{aligned} \tag{D.2.5}$$

D.3 A Special Version of the Functional Central Limit Theorem and Some Consequences

In Subsection 4.6.1 (see Lemma 4.6.5), we showed that the sequence of normalized queues $\{Q^{(P,n)}\}$ in the pooled system (as defined in (4.6.3)) converges almost surely, uniformly to the process $\bar{Q}^P$ defined in (4.4.2). This result addresses the first-order approximation of the pooled queue length process. We now focus on their second-order approximations. Let us start with the following general auxiliary result.

Lemma D.3.1. *Suppose $\tilde{N}^+$ and $\tilde{N}^-$ are independent unit Poisson processes observed on $[0,T]$ and let $\{\Lambda_n^+\}$ and $\{\Lambda_n^-\}$ be sequences in $\mathcal{A}_+$ such that there exist Λ^+ and Λ^- in $\mathcal{A}_+$ satisfying*

$$\begin{aligned} \sqrt{n}(\Lambda_n^+ - \Lambda^+) &\to 0, \\ \sqrt{n}(\Lambda_n^- - \Lambda^-) &\to 0, \end{aligned} \tag{D.3.1}$$

as $n \to \infty$, in the uniform topology. Moreover, let $\bar{X} = \Lambda^+ - \Lambda^-$ and let

$$\hat{X}^n = \sqrt{n}\left(\frac{1}{n}X^n - \bar{X}\right), \quad \textit{for all } n, \tag{D.3.2}$$

where $X^n = N^+(\Lambda_n^+) - N^-(\Lambda_n^-)$ *for all* n*. Then*

$$\hat{X}^n \Rightarrow \hat{X}, \text{ as } n \to \infty, \tag{D.3.3}$$

in the uniform topology, where $\hat{X} \stackrel{(d)}{=} W(\Lambda^+ + \Lambda^-)$ *for* W *a standard Brownian motion.*

Proof. A combination of Theorem 14.6 and the lemma on p. 151 in [Bil99] with the continuity of addition in the uniform topology yields

$$\sqrt{n}\left(\frac{1}{n}X^n - (\Lambda_n^+ - \Lambda_n^-)\right) \Rightarrow \hat{X},$$

uniformly. Invoking assumption (D.3.1) and Theorem 3.9 of [Bil99], we complete the proof. □

The next lemma is dedicated to the queue lengths generated by the netput processes from the previous one.

Lemma D.3.2. *In addition to all the assumptions and the notation from Lemma D.3.1, we set* $\hat{Q}^n = \sqrt{n}(\frac{1}{n}Q^n - \bar{Q})$*, where*

$$Q^n = \Gamma(X^n) \tag{D.3.4}$$

and

$$\bar{Q} = \Gamma(\bar{X}). \tag{D.3.5}$$

Then as $n \to \infty$

$$\hat{Q}_t^n \Rightarrow \hat{Q}_t, \text{for every } t \in [0, T], \tag{D.3.6}$$

where

$$\hat{Q}_t = \hat{X}_t + \sup_{s \in \Phi_{-\bar{X}}(t)} [-\hat{X}_s]$$

with

$$\Phi_{-\bar{X}}(t) = \{s \leq t : -\bar{X}_s = \sup_{u \leq t}[-\bar{X}_u]\}.$$

Proof. We start by using the Skorokhod representation theorem (see, e.g., Theorem 6.7. in [Bil99]) to accommodate all the processes involved on a common probability space so that the convergence in (D.3.3) holds in the almost sure sense.

For every n, the process whose limit we look at in (D.3.6) by definition equals

$$\hat{Q}^n = \sqrt{n}\left(\frac{1}{n}Q^n - \bar{Q}\right).$$

Using the defining equalities (D.3.4) and (D.3.5) and the homogeneity of the one-sided reflection map, we transform the last expression into

$$\hat{Q}^n = \sqrt{n}\left(\Gamma\left(\frac{1}{n}X^n\right) - \Gamma(\bar{X})\right).$$

From (D.3.2) we obtain that $\frac{1}{n}X^n = \bar{X} + \frac{1}{\sqrt{n}}\hat{X}^n$. Inserting this result into the last display, we get

$$\hat{Q}^n = \sqrt{n}\left(\Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}^n\right) - \Gamma(\bar{X})\right). \tag{D.3.7}$$

By Lemma 13.5.1 of [Whi02b], the mapping Γ is Lipschitz continuous in the uniform metric on $\mathcal{D}$ with the Lipschitz constant 2. Thus, we have that

$$\begin{aligned}\sqrt{n}\left\|\Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}^n\right) - \Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}\right)\right\| &\leq 2\sqrt{n}\left\|\bar{X} + \frac{1}{\sqrt{n}}\hat{X}^n - \left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}\right)\right\| \\ &= 2\|\hat{X}^n - \hat{X}\|.\end{aligned}$$

Due to Lemma D.3.1 the right-hand side in the above display perishes in the limit in the almost sure sense on the newly constructed probability space created for the purpose of this proof by means of the Skorokhod representation theorem. In particular, for every $t \in [0, T]$, we have

$$\sqrt{n}\left|\Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}^n\right)_t - \Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}\right)_t\right| \to 0, \text{ as } n \to \infty, \tag{D.3.8}$$

with probability 1. On the other hand, by Theorem 9.5.1 of [Whi02b], we have that for every $t \in [0, T]$

$$\sqrt{n}\left(\Gamma\left(\bar{X} + \frac{1}{\sqrt{n}}\hat{X}\right)_t - \Gamma(\bar{X})_t\right) \to \hat{Q}_t, \textit{ a.s.}, \tag{D.3.9}$$

on the new probability space. Combining the limits (D.3.8) and (D.3.9) through the equality (D.3.7), we get the desired result. □

We intend to employ the last result through the following corollary.

Corollary D.3.3. *Define the sequence of processes* $\hat{Q}^{(P,n)} = \sqrt{n}(\frac{1}{n}Q^{(P,n)} - \bar{Q}^P)$. *Then the following convergence in distribution holds true for every* $t \in [0, T]$*:*

$$\hat{Q}_t^{(P,n)} \Rightarrow \hat{Q}_t^P, \tag{D.3.10}$$

where

$$\hat{Q}_t^P = \hat{X}_t^P + \sup_{s \in \Phi_{-\bar{X}^P}(t)} [-\hat{X}_s^P] \tag{D.3.11}$$

with

$$\Phi_{-\bar{X}^P}(t) = \{s \le t : -\bar{X}^P_s = \sup_{u \le t}[-\bar{X}^P_u]\}$$

while $\bar{X}^P = \mathcal{I}(\lambda - \mu^2)$ *and* $\hat{X}^P \overset{(d)}{=} W(\mathcal{I}(\lambda + \mu^2))$ *for* W *a standard Brownian motion.*

Before exhibiting the last result of this section, we digress briefly to establish the following general lemma.

Lemma D.3.4. *If the random variable* X *is normally distributed and* Y *is another random variable on the same probability space and independent from* X*, then* $\mathbb{P}[X + Y = \alpha] = 0$*, for every real* α*.*

Proof. This is a matter of simple computation in which we denote the characteristic functions of X, Y and $X + Y$ by ϕ_X, ϕ_Y and ϕ_{X+Y}, respectively. By independence of X and Y, we have that $\phi_{X+Y} = \phi_X \phi_Y$. Therefore, for every real ξ,

$$|\phi_{X+Y}(\xi)| = |\phi_X(\xi)||\phi_Y(\xi)| \le |\phi_X(\xi)|.$$

Hence,

$$\int_{-\infty}^{\infty} |\phi_{X+Y}(\xi)| \, d\xi \le \int_{-\infty}^{\infty} |\phi_X(\xi)| \, d\xi < \infty.$$

We conclude that ϕ_{X+Y} is integrable and, therefore, that the random variable $X + Y$ admits a density. The posited claim of the lemma is a simple consequence of this fact. ☐

Corollary D.3.5. *Let* $\hat{Q}^P$ *be as defined in (D.3.11). Then, for every* $t \in [0, T]$ *such that* $\bar{Q}^P_t > 0$*, we have that*

$$\mathbb{P}[\hat{Q}^{(P,n)}_t \in (\alpha, \beta]] \to \mathbb{P}[\hat{Q}^P_t \in (\alpha, \beta]],$$

for every $\alpha, \beta \in \mathbb{R}$*.*

Proof. For every t such that $\bar{Q}^P_t > 0$, defining $\bar{t} = \sup \Phi_{-\bar{X}^P}(t)$ and using the fact that $\hat{X}^P$ is a time-changed Brownian motion, we have

$$\mathbb{P}[\hat{Q}^P_t = \alpha] = \mathbb{P}[\hat{Q}^P_t - \hat{Q}^P_{\bar{t}} + \hat{Q}^P_{\bar{t}} = \alpha] = \mathbb{P}[W(\mathcal{I}_t(\lambda + \mu^2)) - W(\mathcal{I}_{\bar{t}}(\lambda + \mu^2)) + \hat{Q}^P_{\bar{t}} = \alpha].$$

Moreover, the definition of $\hat{Q}^P$ leads us to conclude that the process $\hat{Q}^P$ is itself Markov and we get that the random variable $W(\mathcal{I}_t(\lambda + \mu^2)) - W(\mathcal{I}_{\bar{t}}(\lambda + \mu^2))$ is normally distributed and independent of $\hat{Q}^P_{\bar{t}}$. Hence, Lemma D.3.4 and the last display yield

$$\mathbb{P}[\hat{Q}^P_t = \alpha] = 0. \tag{D.3.12}$$

In words, this means that the random variable $\hat{Q}^P_t$ admits a density. By the same token, we get also $\mathbb{P}[\hat{Q}^P_t = \beta] = 0$. In other words, the random variable $\hat{Q}^P_t$ does not charge the boundary of the set (α, β), so by Portmanteau's theorem and claim (D.3.10) of Corollary D.3.3, we have that $\mathbb{P}[\hat{Q}^{(P,n)}_t \in (\alpha, \beta)] \to \mathbb{P}[\hat{Q}^P_t \in (\alpha, \beta)]$ as $n \to \infty$. ☐

D.4 Auxiliary Stochastic Control results

Here we prove that the probability of the sum of the queue lengths in tandem queue (when the admissible sequence $\{\mu_n\}$ defined in (4.6.9) is used) being strictly greater than the length of the pooled queue vanishes in the limit. Since we are concentrating solely on the performance of the specific admissible sequence $\{\mu_n\}$, we suppress from the notation the dependence of the stochastic processes involved (queue lengths, netputs, etc.) on the service discipline. With this convention, symbolically the result we aim to prove reads as

$$\mathbb{P}[\exists\, t \in [0,T] : Q_t^{(1,n)} + Q_t^{(2,n)} \neq Q_t^{(P,n)}] \to 0, \text{ as } n \to \infty. \tag{D.4.1}$$

Invoking Lemma 4.6.8, we conclude that the desired claim (D.4.1) is equivalent to

$$\mathbb{P}[\exists\, t \in [0,T] : Q_t^{(1,n)} + Q_t^{(2,n)} > Q_t^{(P,n)}] \to 0, \text{ as } n \to \infty. \tag{D.4.2}$$

Some simple algebra (similar to the calculations of the proof of Lemma 4.6.8), yields that the statement (D.4.2) is, in turn, equivalent to

$$\mathbb{P}[\exists\, t \in [0,T] : L_t^{(2,n)} > L_t^{(P,n)}] \to 0, \text{ as } n \to \infty. \tag{D.4.3}$$

For all n, we define the following subset of the set $S^{(n)}$ introduced in (4.6.10):

$$R^{(n)} = \{(q_1, q_2) \in S^{(n)} : q_1 > 0 \,\text{and}\, q_2 = 0\}.$$

Next, we construct three sequences of random times in the following way

$$\begin{aligned}
\xi_1^{(n)} &= \inf\{t > 0 : (Q_t^{(1,n)}, Q_t^{(2,n)}) \in nR^{(n)}\} \wedge T;\\
\chi_i^{(n)} &= \inf\{t > \xi_i^{(n)} : (Q_t^{(1,n)}, Q_t^{(2,n)}) \notin nR^{(n)}\} \wedge T; \text{ for } i \geq 1;\\
\eta_i^{(n)} &= \inf\{t > \xi_i^{(n)} : N_2^-(n\mathcal{I}_t(\mu^2)) > N_2^-(n\mathcal{I}_{\xi_i^{(n)}}(\mu^2))\} \wedge T; \text{ for } i \geq 1;\\
\xi_i^{(n)} &= \inf\{t > \chi_{i-1}^{(n)} : (Q_t^{(1,n)}, Q_t^{(2,n)}) \in nR^{(n)}\} \wedge T; \text{ for } i > 1.
\end{aligned}$$

It is obvious from their definitions that all three sequences of random times are, in fact, sequences of stopping times with respect to the filtration $\{\mathcal{H}_t^{(n)}\}$ constructed in Appendix B.3. Also, since the driving Poisson processes N_1^+, N_1^- and N_2^- are assumed to be independent, for every i the random variables $\chi_i^{(n)}$ and $\eta_i^{(n)}$ are also independent.

A sequence of events of particular interest to us at the moment are the following

$$\{\exists\, i : \chi_i^{(n)} > \eta_i^{(n)}\}. \tag{D.4.4}$$

Let us describe the rationale behind our focus on these events. At each stopping time $\xi_i^{(n)}$, the tandem system is in such a position that the second queue is empty. The uncontrolled service process in the second station $N_2^-(n\mathcal{I}(\mu^2))$ may meanwhile be serving at its given rate. In the case that its exponential service is completed earlier than the service period of the controlled process $N_1^-(n\mathcal{I}(\mu_n))$, the lower regulator in the second queue will increase, and that increase will not be matched (as yet, at least) by the regulator in the pooled queue (as the pooled buffer contains at least the content of the buffer of the first queue at that time). Looking at the definition of the sequence $\{\chi_i^{(n)}\}$, we see that stating that the pair of queue lengths in the tandem queue exits the region $R^{(n)}$ at the time $\chi_i^{(n)}$ precisely means that a service in the first station was completed at the time $\chi_i^{(n)}$. The probability of the event in expression (D.4.4), hence, dominates the probability from (D.4.3). Thus, our course of action is to prove that

$$\mathbb{P}[\exists\, i : \chi_i^{(n)} > \eta_i^{(n)}] \to 0, \text{ as } n \to \infty. \tag{D.4.5}$$

For every n, let the set $A^{(n)}$ be defined as in (4.6.11). Then the probability in the display (D.4.5) can be bounded from above as follows

$$\begin{aligned}\mathbb{P}[\exists\, i : \chi_i^{(n)} > \eta_i^{(n)}] &= \mathbb{P}[\{\exists\, i : \chi_i^{(n)} > \eta_i^{(n)}\} \cap A^{(n)}] + \mathbb{P}[\{\exists\, i : \chi_i^{(n)} > \eta_i^{(n)}\} \cap (A^{(n)})^c] \\ &\leq \mathbb{P}[\{\exists\, i \leq n^{1+\epsilon} : \chi_i^{(n)} > \eta_i^{(n)}\} \cap A^{(n)}] + \mathbb{P}[(A^{(n)})^c] \\ &\leq \mathbb{P}[\exists\, i \leq n^{1+\epsilon} : \chi_i^{(n)} > \eta_i^{(n)}] + \mathbb{P}[(A^{(n)})^c].\end{aligned} \tag{D.4.6}$$

The first probability in the final line of (D.4.6) can be rewritten as

$$\mathbb{P}[\exists\, i \leq n^{1+\epsilon} : \chi_i^{(n)} > \eta_i^{(n)}] = 1 - \mathbb{P}[\forall\, i \leq n^{1+\epsilon} : \chi_i^{(n)} \leq \eta_i^{(n)}]. \tag{D.4.7}$$

For every n, we can interpret the pairs of random variables $\chi_i^{(n)}$ and $\eta_i^{(n)}$ for $i \leq n^{1+\epsilon}$ as the times of the first jumps in a pair of Poisson processes. Due to the strong Markov property of the state process $(Q^{(1,n)}, Q^{(2,n)})$, for every n, we can assume that the pairs of Poisson processes are independent. Along with Assumption 4.6.10, the above observations imply that Corollary D.1.3 is applicable to the probabilities stated in to the right-hand side of expression (D.4.7). Namely, we get that

$$\mathbb{P}[\exists\, i \leq n^{1+\epsilon} : \chi_i^{(n)} > \eta_i^{(n)}] \to 0, \text{ as } n \to \infty. \tag{D.4.8}$$

Combining (D.4.8) and Lemma 4.6.13 with expression (D.4.6), we obtain the limit announced in (D.4.5).

To sum up, we have just proven the following result.

Lemma D.4.1.

$$\mathbb{P}[\exists\, t \in [0,T] : Q_t^{(1,n)} + Q_t^{(2,n)} \neq Q_t^{(P,n)}] \to 0, \textit{ as } n \to \infty.$$

D.5 Asymptotic Optimality - Finite Buffers: Proofs.

Lemma D.5.1. *Let a, b and c be nonnegative real numbers. Then the following inequality holds true.*

$$a \vee b - a \vee c \leq |b - c|.$$

Proof. Let us denote by d the term on the left-hand side of the posited inequality, i.e., let $d = a \vee b - a \vee c$. We intend to simply exhaust all the possible cases, with respect to the ordering of a, b and c.

- If $a \leq b \wedge c$, then $d = b - c \leq |b - c|$.
- If $a \geq b \vee c$, then $d = 0 \leq |b - c|$.
- If $b \leq a \leq c$, then $d = a - c \leq 0 \leq |b - c|$.
- If $c \leq a \leq b$, then $d = b - a \leq b - c \leq |b - c|$.

□

The following is a restatement of Lemma 4.8.1.

Lemma D.5.2. *Consider an admissible service discipline $\mu \in \mathcal{B}$. Then there exists a service discipline $\tilde{\mu} \in \mathcal{B}$ such that the following relations hold almost surely:*

1. $L^{(1)}(\tilde{\mu}) \equiv 0$;
2. $U^{(2)}(\tilde{\mu}) \equiv 0$;
3. $U^{(1)}(\tilde{\mu}) \leq U^{(1)}(\mu) + U^{(2)}(\mu)$.

Proof. It is useful to introduce the following sequence of random times of importance for the tandem queue over the period $[0, T]$, i.e., the times of arrivals to the first station and times of completion for services in the second station. Formally, we set

$$\begin{aligned} &T_0 = 0, \\ &T_i = \inf\{t > T_{i-1} : \Delta N_1^+(\mathcal{I}_t(\lambda)) + \Delta N_1^-(\mathcal{I}_t(\mu)) + \Delta N_2^-(\mathcal{I}_t(\mu^2)) > 0\}, \end{aligned} \tag{D.5.1}$$

for every $i \geq 1$. In words, the sequence $\{T_i\}$ records the times of arrivals and times of potential departures from the first or the second station. Clearly, the random times introduced in the last display are stopping times with respect to the filtration $\{\mathcal{H}_t^{(n)}\}$. Moreover, by independence of the driving Poisson processes, it is almost sure that two events of importance never occur simultaneously.

Next, we tag two special subsequences of the sequence $\{T_i\}$.

$$F_0 = 1,$$
$$F_i = \begin{cases} 1 & \text{if } Q^{(1)}_{T_{i-1}} - \Delta N_1^-(\mathcal{I}_{T_i}(\mu)) \leq 0 \\ 0 & \text{otherwise} \end{cases} \quad \text{for every } i \geq 1,$$
$$G_0 = 0,$$
$$G_i = \begin{cases} 1 & \text{if } Q^{(2)}_{T_{i-1}} + \Delta N_1^-(\mathcal{I}_{T_i}(\mu)) \geq K_2 \\ 0 & \text{otherwise} \end{cases} \quad \text{for every } i \geq 1.$$

The sequence $\{F_i\}$ accounts for all the stopping times in the sequence $\{T_i\}$ such that at that time the first queue either sank to zero, or there was need to regulate the first queue to keep it nonnegative. At the same time, the sequence $\{G_i\}$ flags all the times among $\{T_i\}$ such that the second queue either gets to the level K_2, or there is upper regulation necessary due to the buffer capacity in the second station.

For all $t \in [0, T]$, we define a random index $i(t) = \sup\{i \geq 0 : T_i \leq t\}$. That is, $i(t)$ is the unique integer such that $t \in [T_i, T_{i+1})$. Next, we set, for every $t \in [T_i, T_{i+1})$,

$$\tilde{\mu}_t = \begin{cases} 0 & \text{if } F_{i(t)} = 1 \text{ or } G_{i(t)} = 1 \\ \mu_t & \text{if } F_{i(t)} = G_{i(t)} = 0. \end{cases} \tag{D.5.2}$$

We still need to prove that $\tilde{\mu}$ satisfies all the announced properties.

Admissibility. It is obvious from the construction that the control $\tilde{\mu}$ depends in a non-anticipating manner only on the processes adapted to the filtration $\{\mathcal{H}_t^{(n)}\}$. Thus, it is itself $\{\mathcal{H}_t^{(n)}\}$-predictable.

The constraint placed on the total amount of service available is also satisfied, as the service discipline $\tilde{\mu}$ at all times either coincides with μ or vanishes. The function μ is itself admissible and, thus, almost surely integrates to at most m.

Absence of Lower Regulation in the First Queue. Due to the first defining condition in the first line of (D.5.2), the service $\tilde{\mu}$ "shuts-off" whenever the first queue is at zero. Therefore there is never any need for lower regulation in the first queue.

Absence of Upper Regulation in the Second Queue. The condition including the sequence of tags $\{G_i\}$ in the first case in definition (D.5.2) sets the service discipline $\tilde{\mu}$ to zero at all times when the second queue is at its threshold. This impedes any arrivals into the second station, and thus, prevents all upper regulation in the second station.

Performance. For simplicity's sake, we assume throughout the rest of this proof that the buffer capacities K_1 and K_2 are both integer-valued. The general case involves one small change in the instants at which these capacities are first reached.

We give an inductive proof of the following claims:

$$\begin{aligned}
&Q^{(1)}_t(\mu) = Q^{(1)}_t(\tilde{\mu}), \text{ for every } t \in [T_{i-1}, T_i) \text{ and every } i \in \mathbb{N},\\
&U^{(1)}_t(\tilde{\mu}) \le U^{(1)}_t(\mu) + U^{(2)}_t(\mu), \text{ for every } t \in [T_{i-1}, T_i) \text{ and every } i \in \mathbb{N},\\
&Q^{(2)}_t(\mu) = Q^{(2)}_t(\tilde{\mu}), \text{ for every } t \in [T_{i-1}, T_i) \text{ and every } i \in \mathbb{N}.
\end{aligned} \tag{D.5.3}$$

For $i = 1$, the claims in the above display read

$$\begin{aligned}
&Q^{(1)}_t(\mu) \le Q^{(1)}_t(\tilde{\mu}), \text{ for every } t \in [0, T_1),\\
&U^{(1)}_t(\tilde{\mu}) \le U^{(1)}_t(\mu) + U^{(2)}_t(\mu), \text{ for every } t \in [0, T_1).
\end{aligned} \tag{D.5.4}$$

Since the queues are assumed to be initially empty, for any $t \in [0, T_1)$, both sides of the first line in (D.5.4) are equal to $N^+_1(\mathcal{I}_t(\lambda)) = 0$. Hence, the first claim is trivially satisfied.

Furthermore, clearly, there cannot be any need for upper regulation in either of the queues, as they are both empty until the first arrival into the first station happens. Thus, both sides in the second claim in (D.5.4) are equal to zero.

The length of the second queue is constantly equal to zero on this segment.

Assume that the claim (D.5.3) holds true for all indices smaller than or equal to an index i. Let us prove that the index $i + 1$ necessarily verifies the claim. This will be conducted through an exhaustion of all possible cases of events of importance occurring at time T_i.

Let T_i be such that $\Delta N^+_1(\mathcal{I}_{T_i}(\lambda)) = 1$. The change in the length of the first queue at time T_i can be rewritten as

$$\Delta Q^{(1)}_{T_i}(\mu) = \begin{cases} 1 & \text{if } Q^{(1)}_{T_{i-1}}(\mu) + 1 \le K_1,\\ 0 & \text{if } Q^{(1)}_{T_{i-1}}(\mu) + 1 > K_1. \end{cases} \tag{D.5.5}$$

By the same token

$$\Delta Q^{(1)}_{T_i}(\tilde{\mu}) = \begin{cases} 1 & \text{if } Q^{(1)}_{T_{i-1}}(\tilde{\mu}) + 1 \le K_1,\\ 0 & \text{if } Q^{(1)}_{T_{i-1}}(\tilde{\mu}) + 1 > K_1. \end{cases} \tag{D.5.6}$$

By assumption, $Q^{(1)}_t(\tilde{\mu}) = Q^{(1)}_t(\mu)$, for every $t \le T_i$. Thus, using (D.5.3) and (D.5.6), we conclude that $Q^{(1)}_{T_i}(\tilde{\mu}) = Q^{(1)}_{T_i}(\mu)$, as well.

The exact same argument proves that $\Delta U^{(1)}_{T_i}(\mu) = \Delta U^{(1)}_{T_i}(\tilde{\mu})$. On the other hand, since there are no new arrivals into the second station at time T_i, $\Delta U^{(2)}_{T_i}(\mu) = 0$. The inductive assumption then yields

$$\begin{aligned}
U^{(1)}_{T_i}(\mu) + U^{(2)}_{T_i}(\mu) &= U^{(1)}_{T_i-}(\mu) + U^{(2)}_{T_i}(\mu) + \Delta U^{(1)}_{T_i}(\mu)\\
&\ge U^{(1)}_{T_i-}(\tilde{\mu}) + \Delta U^{(1)}_{T_i}(\tilde{\mu})\\
&= U^{(1)}_{T_i}(\tilde{\mu}).
\end{aligned} \tag{D.5.7}$$

As for the second queue, its length does not change at time T_i as there are no arrivals to this queue or potential departures from it, so the third equality from (D.5.3) carries over directly by the inductive assumption.

In order to be able to extend the claims which we just proved for the instant T_i, to the rest of the interval $[T_i, T_{i+1})$, we need to verify that there are no new potential departures caused by the service $\tilde{\mu}$ during the period $[T_i, T_{i+1})$. However, since $\Delta N_1^+(\mathcal{I}_{T_i}(\lambda)) = 1$ in the present case, we have $Q^{(1)}_{T_i}(\mu) > 0$. By the definition of $\tilde{\mu}$, for each $t \in [T_i, T_{i+1})$, $\tilde{\mu}_t = \mu_t$. We see that the instant of its first jump after time T_i when service rate μ is used coincides with the first jump when $\tilde{\mu}$ is used.

Summing up, there are no changes in any of the primitive processes until time T_{i+1}, so the derived processes do not change either. The claim (D.5.3) is, thus, proven in this case.

In the second case, we assume that $\Delta N_1^-(\mathcal{I}_{T_i}(\mu)) = 1$. The change in the length of the first queue at time T_i is then

$$\begin{aligned}\Delta Q^{(1)}_{T_i}(\mu) &= \begin{cases} -1 & \text{if } Q^{(1)}_{T_{i-1}}(\mu) > 0 \\ 0 & \text{if } Q^{(1)}_{T_{i-1}}(\mu) = 0 \end{cases} \\ &= \begin{cases} -1 & \text{if } F_{i-1} = 0 \\ 0 & \text{if } F_{i-1} = 1. \end{cases}\end{aligned} \tag{D.5.8}$$

Thus, the change in the first queue when the modified service discipline $\tilde{\mu}$ is used equals

$$\Delta Q^{(1)}_{T_i}(\tilde{\mu}) = \begin{cases} \Delta Q^{(1)}_{T_i}(\mu) & \text{if } F_{i-1} = 0 \\ 0 & \text{if } F_{i-1} = 1. \end{cases} \tag{D.5.9}$$

Displays (D.5.8) and (D.5.9), together with the inductive assumption, give us that $Q^{(1)}_{T_i}(\tilde{\mu}) = Q^{(1)}_{T_i}(\mu)$. Clearly, the first queue remains constant over the rest of the random interval $[T_i, T_{i+1})$ when μ is used. As for $Q^{(1)}(\tilde{\mu})$, depending on the value of the tags F_i and G_i, we have the following cases to consider. If $F_i = G_i = 0$, the potential service processes clearly coincide, so that the queues themselves agree, as well. If $F_i = 1$, then queues for both service disciplines vanish over the random segment at hand. If $G_i = 1$ and $F_i = 0$, then the rate μ exceeds the rate $\tilde{\mu}$, which is set to zero. Combining all the stated cases, we get the validity of the claim (D.5.3).

Since there are no new arrivals into the first station at any time in $[T_i, T_{i+1})$, the upper regulator does not increase over that period. As for the second station, when the service discipline μ is used, we have

$$\Delta U^{(2)}_{T_i}(\mu) = \begin{cases} 0 & \text{if } G_{i-1} = 0 \text{ or } F_{i-1} = 1 \\ 1 & \text{if } G_{i-1} = 1 \text{ and } F_{i-1} = 0. \end{cases} \tag{D.5.10}$$

On the other hand, for the effect of the service $\tilde{\mu}$ at the time $\{T_i\}$, we need to first determine the form of $\tilde{\mu}$ in the preceding random interval.

If $F_{i-1} = 1$ or $G_{i-1} = 1$, then $\mu = 0$ throughout $[T_{i-1}, T_i)$, and therefore $\Delta U^{(2)}_{T_i}(\tilde{\mu}) = 0$. Alternatively, if $F_{i-1} = 0$ and $G_{i-1} = 0$, then $\tilde{\mu}$ matches μ over that random interval, and the first times a job is completed coincide for both service disciplines, i.e., $\Delta U^{(2)}_{T_i}(\tilde{\mu}) = \Delta U^{(2)}_{T_i}(\mu) = 0$, by (D.5.10). Simple summation of the kind we have above yields

$$U^{(1)}_{T_i}(\tilde{\mu}) \leq U^{(1)}_{T_i}(\mu) + U^{(2)}_{T_i}(\mu).$$

The result carries over to the rest of the random interval in (D.5.3) in the same fashion as the one above.

The second queue itself has the increment of the following form at time T_i:

$$\Delta Q^{(2)}_{T_i}(\mu) = \begin{cases} 1 & \text{if } F_{i-1} = 0 \text{ and } G_{i-1} = 0 \\ 0 & \text{if } F_{i-1} = 1 \text{ or } G_{i-1} = 1. \end{cases} \tag{D.5.11}$$

On the other hand,

$$\Delta Q^{(2)}_{T_i}(\tilde{\mu}) = \begin{cases} 1 & \text{if } F_{i-1} = 0 \text{ and } G_{i-1} = 0 \\ 0 & \text{if } F_{i-1} = 1 \text{ or } G_{i-1} = 1, \end{cases} \tag{D.5.12}$$

because the service discipline $\tilde{\mu}$ is set to zero in the latter case in the above display. Therefore, $Q^{(2)}_{T_i}(\tilde{\mu}) = Q^{(2)}_{T_i}(\mu)$. A reiteration of the arguments used above for other processes extends this equality to hold for every $t < T_{i+1}$.

Finally, there is a possibility that the event of importance T_i was triggered by a service completion in the second station. Here, the first claim in (D.5.5) is trivially satisfied, as the first queue remains unaltered when either service discipline is utilized.

There is no change in the upper regulators in the first queue or the second queue, as there are no new arrivals into the either queue. This is seen easily using the same rationale as above, which rests on the construction of the controlled service process.

The second queue is changed at time T_i by a (potential) jump down, regardless of whether μ or $\tilde{\mu}$ is used, so the equality in (D.5.5) is satisfied at time T_i. As for the rest of the random interval $[T_i, T_{i+1})$, since there is a jump down at time T_i, we have $G_i = 0$. If $F_i = 0$, then the service processes coincide, and the inductive claim holds true. In case that $F_i = 1$, we set $\tilde{\mu}$ to zero until time T_{i+1}. At the same time, there are no jumps in the service depending on μ by the definition of $\{T_i\}$. □

The following is a restatement of Lemma 4.8.2.

Lemma D.5.3. *For all $\mu \in \mathcal{B}^*$, we have that*

(i) $L^P \leq L^{(2)}(\mu)$, *almost surely;*

(ii) $U^{(1)}(\mu) \geq U^P$, *almost surely.*

Proof. Following the template of the proof of Lemma 4.5.2, we define the two sequences of stopping times follows

$$\begin{aligned}\chi_1 &= \inf\{t > 0 : Q^P_t = K\} \wedge T,\\ \xi_i &= \inf\{t > \chi_i : Q^P_t = 0\} \wedge T, \text{ for } i \geq 1,\\ \chi_i &= \inf\{t > \xi_{i-1} : Q^P_t = K\} \wedge T, \text{ for } i > 1.\end{aligned} \tag{D.5.13}$$

We fix an arbitrary admissible $\mu \in \mathcal{B}^*$ and conclude that, by definition,

$$N_1^+(\mathcal{I}(\lambda)) \geq N_1^-(\mathcal{I}(\mu)) + U^{(1)}(\mu) \geq N_1^-(\mathcal{I}(\mu)). \tag{D.5.14}$$

Again, we intend to use the principle of mathematical induction. On the segment $[0, \chi_1]$, the length of the pooled queue is

$$Q^P = X^P + L^P.$$

The lower regulator L^P can be written for all $t \in [0, \chi_1]$ as

$$L^P_t = \sup_{s \leq t}[-N_1^+(\mathcal{I}_s(\lambda)) + N_2^-(\mathcal{I}_s(\mu^2))].$$

Thanks to the inequality (D.5.14), we obtain

$$L^P_t \leq \sup_{s \leq t}[-N_1^-(\mathcal{I}(\mu)) + N_2^-(\mathcal{I}_s(\mu^2))] = L^{(2)}_t(\mu).$$

Simultaneously, for every $t \leq \chi_1$, we have that $U^P_t = 0$. Due to the nonnegativity of the regulator maps, the second announced inequality holds.

Next, we consider the segment $[\chi_1, \xi_1]$. There is no lower regulation of the pooled queue in this region, so we have that for every $t \in [\chi_1, \xi_1]$,

$$L^P_t = L^P_{\chi_1} \tag{D.5.15}$$

and, therefore,

$$Q^P_t = X^P_t + L^P_{\chi_1} - U^P_t.$$

The regulator maps are by definition nondecreasing, so the equality (D.5.15) and the validity of the announced inequalities on the segment $[0, \chi_1]$ yield

$$L^P_t \leq L^{(2)}_{\chi_1}(\mu) \leq L^{(2)}_t(\mu), \text{ for every } t \in [\chi_1, \xi_1].$$

As for the inequality involving the upper regulators, we have that

$$U^P_t = \sup_{s \leq t}[N_1^+(\mathcal{I}_s(\lambda)) - N_2^-(\mathcal{I}_s(\mu^2)) + L^P_s - K]^+, \text{ for every } t \in [\chi_1, \xi_1].$$

Because the upper regulator U^P increases only when $Q^P = K$, we can rewrite the last equality using (D.5.15) as

$$\begin{aligned}
U_t^P &= \sup_{\chi_1 \le s \le t} [N_1^+(\mathcal{I}_s(\lambda)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_{\chi_1}^P - K]^+ \\
&\le \sup_{\chi_1 \le s \le t} [N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\mu)) - K_1]^+ \\
&\quad + \sup_{\chi_1 \le s \le t} [N_1^-(\mathcal{I}_s(\mu)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_{\chi_1}^P - K_2]^+ \\
&\le \sup_{s \le t} [N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\mu)) - K_1]^+ + \sup_{\chi_1 \le s \le t} [Q_s^{(2)}(\mu) - K_2]^+. \\
&= \sup_{s \le t} [N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\mu)) - K_1]^+ = U_t^{(1)}(\mu), \text{ for every } t \in [\chi_1, \xi_1].
\end{aligned}$$

Let us assume that the posited inequalities hold true on the entire region $[0, \xi_{i-1}]$, for some $i \ge 2$. We will next prove that those claims necessarily carry over from the stated inductive hypothesis to the segment $[\xi_{i-1}, \xi_i]$.

For every $t \in [\xi_{i-1}, \chi_i]$ the pooled queue is strictly below the level K, so there is no need for downward pushing, i.e.,

$$U_t^P = U_{\xi_{i-1}}^P, \text{ for every } t \in [\xi_{i-1}, \chi_i]. \tag{D.5.16}$$

Therefore,

$$Q_t^P = X_t^P + L_t^P - U_{\xi_{i-1}}^P,$$

and

$$L_t^P = L_{\xi_{i-1}}^P \vee \sup_{\xi_{i-1} \le s \le \chi_i} [-N_1^+(\mathcal{I}_s(\lambda)) + N_2^-(\mathcal{I}_s(\mu^2)) + U_{\xi_{i-1}}^P]^+. \tag{D.5.17}$$

From the inductive hypothesis, we conclude that

$$L_{\xi_{i-1}}^P \le L_{\xi_{i-1}}^{(2)}(\mu). \tag{D.5.18}$$

At the same time, using (D.5.14), we get

$$\begin{aligned}
&\sup_{\xi_{i-1} \le s \le \chi_i} [-N_1^+(\mathcal{I}_s(\lambda)) + N_2^-(\mathcal{I}_s(\mu^2)) + U_{\xi_{i-1}}^P]^+ \\
&\qquad \le \sup_{\xi_{i-1} \le s \le \chi_i} [-N_1^-(\mathcal{I}_s(\mu)) - U_s^{(1)}(\mu) + N_2^-(\mathcal{I}_s(\mu^2)) + U_{\xi_{i-1}}^P]^+.
\end{aligned}$$

Since the regulator maps are nondecreasing, using the inductive hypothesis, we get that the quantity on the right-hand side of the last display is smaller than or equal to

$$\sup_{\xi_{i-1} \le s \le \chi_i} [-N_1^-(\mathcal{I}_s(\mu)) + N_2^-(\mathcal{I}_s(\mu^2))]^+.$$

Using the last upper bound in conjunction with (D.5.18) and (D.5.17), we obtain the desired inequality.

The inequality involving the upper regulators is again a simple consequence of the monotonicity of $U^{(1)}$ and (D.5.16).

Finally, we focus on the segment $[\chi_i, \xi_i]$. Here, there is no need for lower regulation in the pooled queue, so

$$L_t^P = L_{\chi_i}^P.$$

We can immediately conclude from the monotonicity of the lower regulator in the second queue that the first proposed inequality holds true in this region.

For every $t \in [\chi_i, \xi_i]$, we have that

$$\begin{aligned}\sup_{\chi_i \le s \le t} &[N_1^+(\mathcal{I}_s(\lambda)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_{\chi_i}^P - K]^+ \\ &\le \sup_{\chi_i \le s \le t}[N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\mu)) - K_1]^+ \\ &\quad + \sup_{\chi_i \le s \le t}[N_1^-(\mathcal{I}_s(\mu)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_{\chi_i}^P - K_2]^+.\end{aligned} \tag{D.5.19}$$

According to the inductive hypothesis, the second term on the right-hand side of the last equation is bounded from above by

$$\sup_{\chi_i \le s \le t}[N_1^-(\mathcal{I}_s(\mu)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_s^{(2)}(\mu) - K_2]^+ = \sup_{\chi_i \le s \le t}[Q_s^{(2)}(\mu) - K_2]^+ = 0.$$

Next, we have that

$$U_t^P = U_{\chi_i}^P \vee \sup_{\chi_i \le s \le t}[N_1^+(\mathcal{I}_s(\lambda)) - N_2^-(\mathcal{I}_s(\mu^2)) + L_{\chi_i}^P - K]^+.$$

Due to (D.5.19), the validity of the first proposed claim on the segment $[\xi_{i-1}, \chi_i]$, and the last equality, we have

$$\begin{aligned}U_t^P &\le U_{\chi_i}^{(1)}(\mu) \vee \sup_{\chi_i \le s \le t}[N_1^+(\mathcal{I}_s(\lambda)) - N_1^-(\mathcal{I}_s(\mu)) - K_1]^+ \\ &= U_t^{(1)}(\mu), \text{ for every } t \in [\chi_i, \xi_i].\end{aligned}$$

□

Bibliography

[AHS05] Bariş Ata, J. M. Harrison, and L. A. Shepp, *Drift rate control of a Brownian processing system*, Ann. Appl. Probab. **15** (2005), no. 2, 1145–1160. MR MR2134100 (2005k:60265)

[AMR04] Rami Atar, Avi Mandelbaum, and Martin I. Reiman, *A Brownian control problem for a simple queueing system in the Halfin-Whitt regime*, Systems Control Lett. **51** (2004), no. 3-4, 269–275.

[Bil99] Patrick Billingsley, *Convergence of probability measures*, second ed., Wiley Series in Probability and Statistics: Probability and Statistics, John Wiley & Sons Inc., New York, 1999, A Wiley-Interscience Publication.

[BKP04] Nikhil Bansal, Tracy Kimbrel, and Kirk Pruhs, *Dynamic speed scaling to manage energy and temperature*, FOCS 2004, 2004, pp. 520–529.

[BW92] Arthur W. Berger and Ward Whitt, *The impact of a job buffer in a token-bank rate-control throttle*, Comm. Statist. Stochastic Models **8** (1992), no. 4, 685–717.

[CADX04] Junxia Chang, Hayriye Ayhan, J. G. Dai, and Cathy H. Xia, *Dynamic scheduling of a multiclass fluid model with transient overload*, Queueing Syst. **48** (2004), no. 3-4, 263–307. MR MR2104107 (2005h:90053)

[CH93] Miklós Csörgő and Lajos 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.

[CR81] 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.

[DR00] Paul Dupuis and Kavita Ramanan, *An explicit formula for the solution of certain optimal control problems on domains with corners*, Teor. Ĭmovīr. Mat. Stat. (2000), no. 63, 32–48. MR MR1870773 (2002h:49045)

[Dup03] Paul Dupuis, *Explicit solution to a robust queueing control problem*, SIAM J. Control Optim. **42** (2003), no. 5, 1854–1875 (electronic). MR MR2046389 (2005b:62171)

[EMW93a] Stephen G. Eick, William A. Massey, and Ward Whitt, *$M_t/G/\infty$ queues with sinusoidal rates*, Man. Sci. **39** (1993), 241–252.

[EMW93b] ______, *The physics of the $M_t/G/\infty$ queue*, Oper. Res. **41** (1993), no. 4, 731–742.

[GK95] Linda V. Green and Peter J. Kolesar, *On the accuracy of the simple peak hour approximation for markovian queues*, Man. Sci. **41** (1995), 1353–1370.

[Hal91] Randolph W. Hall, *Queueing methods for services and manufacturing*, Prentice Hall, NY, 1991.

[Har90] J. Michael Harrison, *Brownian motion and stochastic flow systems*, Robert E. Krieger Publishing Co. Inc., Malabar, FL, 1990, Reprint of the 1985 original. MR MR1092214 (92d:60084)

[HHBM06] Robert Hampshire, Mor Harchol-Balter, and William Massey, *Fluid and diffusion limits for transient sojourn times of processor sharing queues with time varying rates*, Queueing Systems and Their Applications **53** (2006), no. 1/2.

[Hor92] Lajos Horváth, *Strong approximations of open queueing networks*, Math. Oper. Res. **17** (1992), no. 2, 487–508. MR MR1161166 (93d:60055)

[HVM97] J. Michael Harrison and Jan A. Van Mieghem, *Dynamic control of Brownian networks: state space collapse and equivalent workload formulations*, Ann. Appl. Probab. **7** (1997), no. 3, 747–771. MR MR1459269 (98b:60160)

[Kel82] Joseph B. Keller, *Time-dependent queues*, SIAM Rev. **24** (1982), no. 4, 401–412.

[KLRS06] Lukasz Kruk, John Lehoczky, Kavita Ramanan, and Steven Shreve, *An explicit formula for the Skorokhod map on* $[0, a]$, Ann. Probab. (to appear) (2006).

[Kus01] Harold J. Kushner, *Heavy traffic analysis of controlled queueing and communication networks*, Applications of Mathematics (New York), vol. 47, Springer-Verlag, New York, 2001, Stochastic Modelling and Applied Probability. MR MR1834938 (2002c:90003)

[Mas02] William A. Massey, *The analysis of queues with time-varying rates for telecommunications models*, Telecommunication Systems **21** (2002), no. 2-4, 173–204.

[Mey01] Sean P. Meyn, *Sequencing and routing in multiclass queueing networks. I. Feedback regulation*, SIAM J. Control Optim. **40** (2001), no. 3, 741–776 (electronic). MR MR1871453 (2002k:90050)

[Mey03] ______, *Sequencing and routing in multiclass queueing networks. II. Workload relaxations*, SIAM J. Control Optim. **42** (2003), no. 1, 178–217 (electronic). MR MR1982741 (2004e:90024)

[MM95] Avi Mandelbaum and William A. Massey, *Strong approximations for time-dependent queues*, Math. Oper. Res. **20** (1995), no. 1, 33–64.

[MR06] Avishai Mandelbaum and Kavita Ramanan, *Directional derivatives of oblique reflection maps*, Mathematics of Operations Research (to appear) (2006).

[MW93] William A. Massey and Ward Whitt, *Networks of infinite-server queues with nonstationary Poisson input*, Queueing Systems Theory Appl. **13** (1993), no. 1-3, 183–250.

[New68a] G. F. Newell, *Queues with time-dependent arrival rates. I. The transition through saturation*, J. Appl. Probability **5** (1968), 436–451.

[New68b] ______, *Queues with time-dependent arrival rates. II. The maximum queue and the return to equilibrium*, J. Appl. Probability **5** (1968), 579–590.

[New68c] ______, *Queues with time-dependent arrival rates. III. A mild rush hour*, J. Appl. Probability **5** (1968), 591–606.

[New71] ______, *Applications of queueing theory*, Chapman and Hall Ltd., London, 1971, Monographs on Applied Probability and Statistics.

[Sko61] A. V. Skorohod, *Stochastic equations for diffusion processes with a boundary*, Teor. Verojatnost. i Primenen. **6** (1961), 287–298.

[Whi02a] Ward 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.

[Whi02b] ______, *Stochastic-process limits*, Springer Series in Operations Research, Springer-Verlag, New York, 2002, An introduction to stochastic-process limits and their application to queues.

[Wil00] R. J. Williams, *On dynamic scheduling of a parallel server system with complete resource pooling*, Analysis of communication networks: call centres, traffic and performance (Toronto, ON, 1998), Fields Inst. Commun., vol. 28, Amer. Math. Soc., Providence, RI, 2000, pp. 49–71. MR MR1788708 (2001h:90037)