Infinite Horizon Asymptotic Average Optimality for Large-Scale Parallel Server Networks

Ari Arapostathis[a], Guodong Pang[b,*]

[a]*Department of Electrical and Computer Engineering, The University of Texas at Austin, 2501 Speedway, EER 7.824, Austin, TX 78712*
[b]*The Harold and Inge Marcus Dept. of Industrial and Manufacturing Eng., College of Engineering, Pennsylvania State University, University Park, PA 16802*

Abstract

We study infinite-horizon asymptotic average optimality for parallel server networks with multiple classes of jobs and multiple server pools in the Halfin–Whitt regime. Three control formulations are considered: 1) minimizing the queueing and idleness cost, 2) minimizing the queueing cost under a constraints on idleness at each server pool, and 3) fairly allocating the idle servers among different server pools. For the third problem, we consider a class of *bounded-queue, bounded-state* (BQBS) stable networks, in which any moment of the state is bounded by that of the queue only (for both the limiting diffusion and diffusion-scaled state processes). We show that the optimal values for the diffusion-scaled state processes converge to the corresponding values of the ergodic control problems for the limiting diffusion. We present a family of state-dependent Markov balanced saturation policies (BSPs) that stabilize the controlled diffusion-scaled state processes. It is shown that under these policies, the diffusion-scaled state process is exponentially ergodic, provided that at least one class of jobs has a positive abandonment rate. We also establish useful moment bounds, and study the ergodic properties of the diffusion-scaled state processes, which play a crucial role in proving the asymptotic optimality.

Keywords: multiclass multi-pool Markovian queues, Halfin–Whitt (QED) regime, ergodic control (with constraints), fairness, exponential stability, balanced saturation policy (BSP), bounded-queue bounded-state (BQBS) stable networks, asymptotic optimality
2010 MSC: Primary: 60K25, 68M20, 90B22, 90B36

1. Introduction

Large-scale parallel server networks are used to model various service, manufacturing and telecommunications systems; see, e.g., [1–13]. We consider multiclass multi-pool networks operating in the Halfin–Whitt (H–W) regime, where the demand of each class and the numbers of servers in each pool get large simultaneously in an appropriate manner so that the system becomes critically loaded while the service and abandonment rates are fixed. We study optimal control problems of such networks under the infinite-horizon expected average (ergodic) cost criteria, since steady-state performance measures are among the most important metrics to understand the system dynamics. Specifically, we consider the following unconstrained and constrained ergodic control problems (see Sections 3.1 and 4.2): (P1) minimizing the queueing and idleness cost, (P2) minimizing the queueing cost while imposing a dynamic constraint on the idleness of each server pool (e.g., requiring that the long-run average idleness does not

*Corresponding author
Email addresses: ari@ece.utexas.edu (Ari Arapostathis), gup3@psu.edu (Guodong Pang)

Preprint submitted to Elsevier *March 2, 2018*

exceed a given threshold), and (P3) minimizing the queueing cost while requiring fairness on idleness (e.g., the average idleness of each server pool is a fixed proportion of the total average idleness of all server pools). The scheduling policy determines the allocation of service capacity to each class at each time. We consider only work conserving scheduling policies that are non-anticipative and preemptive.

In [14] and [15], we have studied the corresponding ergodic control problems (P1′)–(P2′) for the limiting diffusions arising from such networks (see Section 3.2). Problem (P1′) for multiclass networks ("V" networks) was studied in [14], where a comprehensive study of the ergodic control problem for a broader class of diffusions as well as asymptotic optimality results for the network model were presented. In [15] we have shown that problem (P1′) and (P2′) are well-posed for multiclass multi-pool networks, and presented a full characterization of optimality for the limiting diffusion. We also provided important insights on the stabilizability[1] of the controlled diffusion by employing a leaf elimination algorithm, which is used to derive an explicit expression for the drift. We addressed problem (P3) for the 'N' network model, where we studied its well-posedness, characterized the optimal solutions, and established asymptotic optimality in [16]. For this particular network topology, the fairness constraint requires that the long-run average idleness of the two server pools satisfies a fixed ratio condition.

In this paper we establish the asymptotic optimality for the ergodic control problems (P1)–(P3). In other words, we show that the optimal values for the diffusion-scaled state processes converge to the corresponding values for the limiting diffusion. The main challenge lies in understanding the recurrence properties of the diffusion-scaled state processes for multiclass multi-pool networks in the H–W regime. Despite the recent studies on stability of multiclass multi-pool networks under certain scheduling policies [17–19], the existing results are not sufficient for our purpose. The difficulty is particularly related to the so-called "joint work conservation" (JWC) condition, which requires that no servers are idling unless all the queues are empty, and which plays a key role in the derivation of the limiting diffusion, and the study of discounted control problems in [20, 21]; see Section 2.2 for a detailed discussion. For the limiting diffusion, the JWC condition holds over the entire state space; however, for the diffusion-scaled state process in the n^{th} system (n is the scaling parameter), it holds only in a bounded subset of the state space. As a consequence, a stabilizing control[2] for the limiting diffusion cannot be directly translated to a scheduling policy for the n^{th} system which stabilizes the diffusion-scaled state process in the H–W regime.

Our first main contribution addresses the above mentioned critical issue of stabilizability of the diffusion-scaled state processes. We have identified a family of stabilizing policies for multiclass multi-pool networks, which we refer to as the "Balanced Saturation Policies" (BSPs) (see Definition 5.1). Such a policy strives to keep the state process for each class 'close' to the corresponding steady state quantity, which is state-dependent and dynamic. The specific stabilizing policy for the 'N' network presented in [16] belongs to this family of BSPs. We show that if the abandonment parameter is positive for at least one class, then the diffusion-scaled state processes are exponentially stable under any BSP (see Proposition 5.1).

In addition to the ergodicity properties, moment bounds are also essential to prove asymptotic optimality. An important implication of the exponential ergodicity property proved in [15, Theorem 4.2] is that the controlled diffusion satisfies a very useful moment bound, see (2.29), namely, that any moment (higher than first order) of the state is controlled by the corresponding moments of the queue and idleness. This moment bound is also shown for the n^{th} system (Proposition 6.1). In studying the moment bounds, we have identified an important class of multiclass multi-pool networks, which we refer to as *bounded-queue, bounded-state* (BQBS) stable networks (Section 4). The limiting diffusion of this class of networks has the

[1] We say that a control (policy) is stabilizing, if it results in a finite value for the optimization criterion.

[2] To avoid confusion, 'control' always refers to a control strategy for the limiting diffusion, while 'policy' refers to a scheduling strategy for the pre-limit model.

following important property. Any moment (higher than first order) of the state is controlled by the corresponding moment of the queue alone (see Proposition 6.2). The class of BQBS stable networks contains many interesting examples, including networks with a single dominant class (see Figure 1), and networks with certain parameter constraints, e.g., service rates that are only pool-dependent. It is worth noting that the ergodic control problem with fairness constraints for the general multiclass multi-pool networks may not be well-posed. However, the problems (P3) and (P3′) are well-posed for the BQBS stable networks. In addition, the problems (P1) and (P1′) which only penalize the queueing cost are also well-posed for the BQBS stable networks.

The proof of asymptotic optimality involves the convergence of the value functions, specifically, establishing the lower and upper bounds (see Theorems 3.1, 3.2 and 4.2). In establishing the lower bound, the arguments are analogous to those for the 'N' network in [16] and the 'V' network in [14]. This involves proving the tightness of the mean empirical measures of the diffusion-scaled state process, controlled under some eventually JWC scheduling policy (see Definition 2.2), and also showing that any limit of these empirical measures is an ergodic occupation measure for the limiting diffusion model (see Lemma 6.1).

The proof of the upper bound is the most challenging. We utilize the following important property which arises from a spatial truncation technique for all three problems (P1′)–(P3′): there exists a continuous precise stationary Markov control $\bar{v}_\epsilon$ which is ϵ-optimal for the limiting diffusion control problem, and under which the diffusion is exponentially ergodic (see Lemma 7.1, Corollary 7.1, and [14, Theorem 4.2]). For the n^{th} system, we construct a concatenated admissible policy in the following manner. In the JWC conservation region, we apply a scheduling policy constructed canonically from the Markov control $\bar{v}_\epsilon$ (see Definition 6.1), while outside the JWC region, we apply a fixed BSP. We also show that that under this concatenated policy the diffusion-scaled state process is exponentially ergodic, and its mean empirical measures converge to the ergodic occupation measure of the limiting diffusion associated with the control $\bar{v}_\epsilon$.

1.1. Literature review

There is an extensive literature on scheduling control of multiclass multi-pool networks in the H–W regime. For the infinite-horizon discounted criterion, Atar [20, 21] first studied the unconstrained scheduling control problem under a set of conditions on the network structure, the system parameters, and the running cost function (Assumptions 2 and 3 in [21]). Atar et al. [22] further investigated simplified models with service rates that either only class-dependent, or pool-dependent. Gurvich and Whitt [23–25] studied queue-and-idleness-ratio controls for multiclass multi-pool networks, by proving a state-space-collapse (SSC) property under suitable conditions on the network structure and system parameters (Theorems 3.1 and 5.1 in [23]). For finite-horizon cost criteria, Dai and Tezcan [26, 27] studied scheduling controls of multiclass multi-pool networks, also by proving an SSC property under certain assumptions.

There has also been a lot of activity on ergodic control of multiclass multi-pool networks in the H–W regime, in addition to [14–16] mentioned earlier. For the inverted 'V' model, Armony [28] has shown that the fastest-server-first policy is asymptotically optimal for minimizing the steady-state expected queue length and waiting time, and Armony and Ward [29] have shown that a threshold policy is asymptotically optimal for minimizing the expected queue length and waiting time subject to a "fairness" constraint on the workload division. For multiclass multi-pool networks, Ward and Armony [30] have studied blind fair routing policies, and used simulations to validate their performance, and compared them with non-blind policies derived from the limiting diffusion control problem. Biswas [31] recently studied a specific multiclass multi-pool network with "help" where each server pool has a dedicated stream of a customer class, and can help with other customer classes only when it has idle servers. For this network model, the control policies may not be work-conserving, and the associated controlled diffusion has a uniform stability property, which is not satisfied for general multiclass multi-pool networks.

This work contributes to the understanding of the stability of multiclass multi-pool networks in the H-W regime. Gamarnik and Stolyar [32] studied the tightness of the stationary distributions of the diffusion-scaled state processes under any work conserving scheduling policy for the 'V' network, while ergodicity properties for the limiting diffusion under constant Markov controls are established in [33, 34]. We refer the reader to [17–19] for the stability analysis of a load balancing scheduling policy, "longest-queue freest-server" (LQFS-LB), and a *leaf activity priority* policy, for multiclass multi-pool networks. For the 'N' network with no abandonment, Stolyar [35] studied the stability of a static priority scheduling policy.

1.2. Organization of the paper

In the subsection which follows we summarize the notation used in the paper. In Section 2.1, we describe the model and the scheduling control problems, and in Section 2.2, we discuss the JWC condition. In Section 2.3, we state some basic properties of the diffusion-scaled processes and the control parameterization, which leads to the diffusion limit. In Section 2.4, we review some relevant properties of the limiting diffusion from [15]. We state the control objectives of the problems (P1) and (P2) in Section 3.1, and the corresponding diffusion control problems (P1$'$) and (P2$'$) in Section 3.2, and summarize the asymptotic optimality results in Section 3.3. In Section 4 we describe the BQBS stable networks and study the fairness problems (P3) and (P3$'$). In Section 5, we introduce the family of stabilizing BSPs, and show that under these we have exponential stability. In Section 6, we focus on the ergodic properties of the n^{th} system, including the moment bounds, convergence of mean empirical measures and a stability preserving property in the JWC region. In Section 7, we complete the proofs of the lower and upper bounds of the three problems. We conclude in Section 8.

1.3. Notation

The symbol $\mathbb{R}$ denotes the field of real numbers, and $\mathbb{R}_+$ and $\mathbb{N}$ denote the sets of nonnegative real numbers and natural numbers, respectively. The minimum (maximum) of two real numbers a and b, is denoted by $a \wedge b$ ($a \vee b$). Define $a^+ := a \vee 0$ and $a^- := -(a \wedge 0)$. The integer part of a real number a is denoted by $\lfloor a \rfloor$. We also let $e := (1, \dots, 1)^{\mathsf{T}}$.

For a set $A \subset \mathbb{R}^d$, we use $\bar{A}$, A^c, and $\mathbb{1}_A$ to denote the closure, the complement, and the indicator function of A, respectively. A ball of radius $r > 0$ in $\mathbb{R}^d$ around a point x is denoted by $B_r(x)$, or simply as B_r if $x = 0$. The Euclidean norm on $\mathbb{R}^d$ is denoted by $|\cdot|$, $x \cdot y$ denotes the inner product of $x, y \in \mathbb{R}^d$, and $\|x\| := \sum_{i=1}^d |x_i|$.

We let $\mathcal{C}_c^\infty(\mathbb{R}^d)$ denote the set of smooth real-valued functions on $\mathbb{R}^d$ with compact support. For a Polish space $\mathcal{X}$, we denote by $\mathcal{P}(\mathcal{X})$ the space of probability measures on the Borel subsets of $\mathcal{X}$ under the Prokhorov topology. For $\nu \in \mathcal{P}(\mathcal{X})$ and a Borel measurable map $f \colon \mathcal{X} \to \mathbb{R}$, we often use the abbreviated notation $\nu(f) := \int_{\mathcal{X}} f \, \mathrm{d}\nu$. The quadratic variation of a square integrable martingale is denoted by $\langle \cdot , \cdot \rangle$. For any path $X(\cdot)$ of a càdlàg process, we use the notation $\Delta X(t)$ to denote the jump at time t.

2. The Model

All random variables introduced below are defined on a complete probability space $(\Omega, \mathfrak{F}, \mathbb{P})$ and $\mathbb{E}$ denotes the associated expectation operator.

2.1. The multiclass multi-pool network model

We consider a sequence of network systems with the associated variables, parameters and processes indexed by n. Each of these, is a multiclass multi-pool Markovian network with I classes of customers and J server pools, labeled as $1, \dots, I$ and $1, \dots, J$, respectively. Let $\mathcal{I} = \{1, \dots, I\}$ and $\mathcal{J} = \{1, \dots, J\}$. Customers of each class form their own queue and are served in the first-come-first-served (FCFS) service discipline. The buffers of all classes are assumed to have infinite capacity. Customers can abandon/renege while waiting in queue. Each class of customers can be served by a subset of server pools, and each server pool can serve a subset of customer classes. We let $\mathcal{J}(i) \subset \mathcal{J}$, denote the subset of server pools that can serve class i customers, and $\mathcal{I}(j) \subset \mathcal{I}$ the subset of customer classes that can be served by server pool j. We form a bipartite graph $\mathcal{G} = (\mathcal{I} \cup \mathcal{J}, \mathcal{E})$ with a set of edges defined by $\mathcal{E} = \{(i,j) \in \mathcal{I} \times \mathcal{J} \colon j \in \mathcal{J}(i)\}$, and use the notation $i \sim j$, if $(i,j) \in \mathcal{E}$, and $i \nsim j$, otherwise. We assume that the graph $\mathcal{G}$ is a tree.

For each $j \in \mathcal{J}$, let N_j^n be the number of servers (statistically identical) in server pool j. Set $N^n = (N_j^n)_{j\in\mathcal{J}}$. Customers of class $i \in \mathcal{I}$ arrive according to a Poisson process with rate $\lambda_i^n > 0$, and have class-dependent exponential abandonment rates $\gamma_i^n \geq 0$. These customers are served at an exponential rate $\mu_{ij}^n > 0$ at server pool j, if $i \sim j$, and we set $\mu_{ij}^n = 0$, if $i \nsim j$. Thus, the set of edges $\mathcal{E}$ can thus be written as $\mathcal{E} = \bigl\{(i,j) \in \mathcal{I} \times \mathcal{J} \,:\, \mu_{ij}^n > 0\bigr\}$. We assume that the customer arrival, service, and abandonment processes of all classes are mutually independent. We define

$$\mathbb{R}_+^{\mathcal{G}} \;\coloneqq\; \bigl\{\xi = [\xi_{ij}] \in \mathbb{R}_+^{I\times J} \colon \xi_{ij} = 0 \;\text{ for } i \nsim j\bigr\}\,,$$

and analogously define $\mathbb{Z}_+^{\mathcal{G}}$.

2.1.1. The Halfin–Whitt regime

We study these multiclass multi-pool networks in the Halfin–Whitt regime (or the Quality-and-Efficiency-Driven (QED) regime), where the arrival rates of each class and the numbers of servers of each server pool grow large as $n \to \infty$ in such a manner that the system becomes critically loaded. Throughout the paper, the set of parameters is assumed to satisfy the following.

Parameter Scaling. There exist positive constants λ_i and ν_j, nonnegative constants γ_i and μ_{ij}, with $\mu_{ij} > 0$ for $i \sim j$ and $\mu_{ij} = 0$ for $i \nsim j$, and constants $\hat\lambda_i$, $\hat\mu_{ij}$ and $\hat\nu_j$, such that the following limits exist as $n \to \infty$.

$$\frac{\lambda_i^n - n\lambda_i}{\sqrt{n}} \;\to\; \hat\lambda_i\,, \qquad \sqrt{n}\,(\mu_{ij}^n - \mu_{ij}) \;\to\; \hat\mu_{ij}\,, \qquad \frac{N_j^n - n\nu_j}{\sqrt{n}} \;\to\; \hat\nu_j\,, \qquad \gamma_i^n \;\to\; \gamma_i\,. \tag{2.1}$$

Fluid scale equilibrium. We assume that the linear program (LP) given by

$$\text{Minimize} \quad \max_{j\in\mathcal{J}} \sum_{i\in\mathcal{I}} \xi_{ij}\,, \quad \text{subject to} \quad \sum_{j\in\mathcal{J}} \mu_{ij}\nu_j\xi_{ij} \;=\; \lambda_i\,,\; i\in\mathcal{I}\,, \quad \text{and} \quad [\xi_{ij}] \in \mathbb{R}_+^{\mathcal{G}}\,,$$

has a unique solution $\xi^* = [\xi_{ij}^*] \in \mathbb{R}_+^{\mathcal{G}}$ satisfying

$$\sum_{i\in\mathcal{I}} \xi_{ij}^* \;=\; 1, \quad \forall j \in \mathcal{J}\,, \quad \text{and} \quad \xi_{ij}^* > 0 \quad \text{for all } i \sim j\,. \tag{2.2}$$

This assumption is referred to as the *complete resource pooling* condition [21, 36]. It implies that the graph $\mathcal{G}$ is a tree [21, 36].

We define $x^* = (x_i^*)_{i\in\mathcal{I}} \in \mathbb{R}_+^I$, and $z^* = [z_{ij}^*] \in \mathbb{R}_+^{\mathcal{G}}$ by

$$x_i^* \;=\; \sum_{j\in\mathcal{J}} \xi_{ij}^*\nu_j\,, \qquad z_{ij}^* = \xi_{ij}^*\nu_j\,. \tag{2.3}$$

The vector x^* can be interpreted as the steady-state total number of customers in each class, and the matrix z^* as the steady-state number of customers in each class receiving service, in the fluid scale. Note that the steady-state queue lengths are all zero in the fluid scale. The quantity ξ_{ij}^* can be interpreted as the steady-state fraction of service allocation of pool j to class-i jobs in the fluid scale. It is evident that (2.2) and (2.3) imply that $e \cdot x^* = e \cdot \nu$, where $\nu \coloneqq (\nu_j)_{j\in\mathcal{J}}$.

2.1.2. The state descriptors

For $i \in \mathcal{I}$, let $X_i^n = \{X_i^n(t) \colon t \ge 0\}$ and $Q_i^n = \{Q_i^n(t) \colon t \ge 0\}$ be the number of class i customers in the system and in the queue, respectively, and for $j \in \mathcal{J}$, let $Y_j^n = \{Y_j^n(t) \colon t \ge 0\}$, be the number of idle servers in pool j. We also let $Z_{ij}^n = \{Z_{ij}^n(t) \colon t \ge 0\}$ denote the number of class i customers being served in server pool j. Set $X^n = (X_i^n)_{i\in\mathcal{I}}$, $Y^n = (Y_j^n)_{j\in\mathcal{J}}$, $Q^n = (Q_i^n)_{i\in\mathcal{I}}$, and $Z^n = (Z_{ij}^n)_{i\in\mathcal{I},\, j\in\mathcal{J}}$. For each $t \ge 0$, we have the fundamental balance equations

$$
\begin{aligned}
X_i^n(t) &= Q_i^n(t) + \sum_{j\in\mathcal{J}(i)} Z_{ij}^n(t) \qquad \forall\, i \in \mathcal{I}\,,\\
N_j^n &= Y_j^n(t) + \sum_{i\in\mathcal{I}(j)} Z_{ij}^n(t) \qquad \forall\, j \in \mathcal{J}\,.
\end{aligned}
\tag{2.4}
$$

2.1.3. Scheduling control

The control process is Z^n. We only consider work conserving scheduling policies that are non-anticipative and preemptive. Work conservation requires that the processes Q^n and Y^n satisfy

$$Q_i^n(t) \wedge Y_j^n(t) \;=\; 0 \qquad \forall i \sim j\,, \quad \forall t \ge 0\,.$$

In other words, whenever there are customers waiting in queues, if a server becomes free and can serve one of the customers, the server cannot idle and must decide which customer to serve and start service immediately. Service preemption is allowed, that is, service of a customer can be interrupted at any time to serve some other customer of another class and resumed at a later time.

For $(x, z) \in \mathbb{Z}_+^I \times \mathbb{Z}_+^{\mathcal{G}}$, we define

$$
\begin{aligned}
q_i(x,z) &:= x_i - \sum_{j\in\mathcal{J}} z_{ij}\,, \quad i \in \mathcal{I}\,,\\
y_j^n(z) &:= N_j^n - \sum_{i\in\mathcal{J}} z_{ij}\,, \quad j \in \mathcal{J}\,,
\end{aligned}
\tag{2.5}
$$

and the *action set* $\mathcal{Z}^n(x)$ by

$$\mathcal{Z}^n(x) \;:=\; \bigl\{z \in \mathbb{Z}_+^{\mathcal{G}} : q_i(x,z) \wedge y_j^n(z) = 0\,,\; q_i(x,z) \ge 0\,,\; y_j^n(z) \ge 0 \quad \forall\,(i,j) \in \mathcal{E}\bigr\}\,.$$

We denote $y_j(x,z) = y_j^n(x,z)$ whenever no confusion occurs.

Let A_i^n, S_{ij}^n, and R_i^n, $(i,j) \in \mathcal{E}$, be mutually independent rate-1 Poisson processes, and also independent of the initial condition $X_i^n(0)$. Define the σ-fields

$$
\begin{aligned}
\mathcal{F}_t^n &:= \sigma\bigl\{X^n(0), \tilde{A}_i^n(t), \tilde{S}_{ij}^n(t), \tilde{R}_i^n(t) \,:\, i \in \mathcal{I},\; j \in \mathcal{J},\; 0 \le s \le t\bigr\} \vee \mathcal{N}\,,\\
\mathcal{G}_t^n &:= \sigma\bigl\{\delta\tilde{A}_i^n(t,r), \delta\tilde{S}_{ij}^n(t,r), \delta\tilde{R}_i^n(t,r) \,:\, i \in \mathcal{I},\; j \in \mathcal{J},\; r \ge 0\bigr\}\,,
\end{aligned}
$$

where $\mathcal{N}$ is the collection of all $\mathbb{P}$-null sets, and

$$
\begin{aligned}
\tilde{A}_i^n(t) &:= A_i^n(\lambda_i^n t), & \delta\tilde{A}_i^n(t,r) &:= \tilde{A}_i^n(t+r) - \tilde{A}_i^n(t)\,,\\
\tilde{S}_{ij}^n(t) &:= S_{ij}^n\left(\mu_{ij}^n \int_0^t Z_{ij}^n(s)\,\mathrm{d}s\right), & \delta\tilde{S}_{ij}^n(t,r) &:= S_{ij}^n\left(\mu_{ij}^n \int_0^t Z_{ij}^n(s)\,\mathrm{d}s + \mu_{ij}^n r\right) - \tilde{S}_{ij}^n(t)\,,\\
\tilde{R}_i^n(t) &:= R_i^n\left(\gamma_i^n \int_0^t Q_i^n(s)\,\mathrm{d}s\right), & \delta\tilde{R}_i^n(t,r) &:= R_i^n\left(\gamma_i^n \int_0^t Q_i^n(s)\,\mathrm{d}s + \gamma_i^n r\right) - \tilde{R}_i^n(t)\,.
\end{aligned}
$$

The filtration $\boldsymbol{\mathcal{F}}^n := \{\mathcal{F}_t^n \colon t \ge 0\}$ represents the information available up to time t, and the filtration $\boldsymbol{\mathcal{G}}^n := \{\mathcal{G}_t^n \colon t \ge 0\}$ contains the information about future increments of the processes.

We say that a scheduling policy Z^n is *admissible* if

(i) $Z^n(t) \in \mathcal{Z}^n(X^n(t))$ a.s. for all $t \geq 0$;

(ii) $Z^n(t)$ is adapted to $\mathfrak{F}^n_t$;

(iii) $\mathfrak{F}^n_t$ is independent of $\mathfrak{G}^n_t$ at each time $t \geq 0$;

(iv) for each $i \in \mathcal{I}$ and $i \in \mathcal{J}$, and for each $t \geq 0$, the process $\delta \tilde{S}^n_{ij}(t, \cdot)$ agrees in law with $S^n_{ij}(\mu^n_{ij}\, \cdot)$, and the process $\delta \tilde{R}^n_i(t, \cdot)$ agrees in law with $R^n_i(\gamma^n_i \cdot)$.

We denote the set of all admissible scheduling policies $(Z^n, \boldsymbol{\mathfrak{F}}^n, \boldsymbol{\mathfrak{G}}^n)$ by $\mathfrak{Z}^n$. Abusing the notation we sometimes denote this as $Z^n \in \mathfrak{Z}^n$. An admissible policy is called stationary Markov if $Z^n(t) = z(X^n(t))$ for some function $z \colon \mathbb{Z}^I_+ \to \mathbb{Z}^{\mathcal{G}}_+$, in which case we identify the policy with the function z.

Under an admissible scheduling policy, the state process X^n can be represented as

$$X^n_i(t) \;=\; X^n_i(0) + A^n_i(\lambda^n_i t) - \sum_{j \in \mathcal{J}(i)} S^n_{ij}\left(\mu^n_{ij} \int_0^t Z^n_{ij}(s)\mathrm{d}s\right) - R^n_i\left(\gamma^n_i \int_0^t Q^n_i(s)\mathrm{d}s\right), \tag{2.6}$$

for $i \in \mathcal{I}$ and $t \geq 0$. Under a stationary Markov policy, X^n is Markov with generator

$$\begin{aligned}\mathcal{L}^z_n f(x) \;&:=\; \sum_{i\in\mathcal{I}} \lambda^n_i \bigl(f(x+e_i) - f(x)\bigr) + \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}(i)} \mu^n_{ij} z_{ij} \bigl(f(x-e_i) - f(x)\bigr)\\ &\qquad + \sum_{i\in\mathcal{I}} \gamma^n_i q_i(x,z) \bigl(f(x-e_i) - f(x)\bigr)\,, \qquad f \in \mathcal{C}(\mathbb{R}^I)\,, \quad x \in \mathbb{Z}^I_+\,. \end{aligned} \tag{2.7}$$

2.2. Joint work conservation

Definition 2.1. We say that an action $z \in \mathcal{Z}^n(x)$ is jointly work conserving (JWC), if

$$e \cdot q(x,z) \wedge e \cdot y^n(z) \;=\; 0\,. \tag{2.8}$$

We define

$$\mathcal{X}^n \;:=\; \bigl\{x \in \mathbb{Z}^I_+ :\;\; e \cdot q(x,z) \wedge e \cdot y^n(x,z) \;=\; 0 \text{ for some } z \in \mathcal{Z}^n(x)\bigr\}\,,$$

with q and y^n defined in (2.5).

Since (2.4) implies that

$$e \cdot (x - N^n) \;=\; e \cdot q(x,z) - e \cdot y^n(z)\,,$$

it is clear that (2.8) is satisfied if and only if

$$e \cdot q(x,z) \;=\; \bigl[e \cdot (x - N^n)\bigr]^+\,, \quad \text{and} \quad e \cdot y^n(z) \;=\; \bigl[e \cdot (x - N^n)\bigr]^-\,.$$

Let

$$\mathit{\Theta}^n(x) \;:=\; \bigl\{(q,y) \in \mathbb{Z}^I_+ \times \mathbb{Z}^J_+ : e \cdot q = [e \cdot (x - N^n)]^+\,,\; e \cdot y = [e \cdot (x - N^n)]^-\bigr\}\,, \qquad x \in \mathbb{Z}^I_+\,.$$

It is evident that the JWC condition can be met at any point $x \in \mathbb{Z}^I_+$ at which the image of $\mathcal{Z}^n(x)$ under the map $z \mapsto \bigl(q(x,z), y^n(z)\bigr)$ defined in (2.5) intersects $\mathit{\Theta}^n(x)$.

Let

$$D_\Psi \;:=\; \bigl\{(\alpha,\beta) \in \mathbb{R}^I \times \mathbb{R}^J \,:\, e \cdot \alpha = e \cdot \beta\bigr\}\,.$$

As shown in Proposition A.2 of [20], provided that $\mathcal{G}$ is a tree, there exists a unique linear map $\Psi = [\Psi_{ij}] \colon D_\Psi \to \mathbb{R}^{I \times J}$ solving

$$\sum_j \Psi_{ij}(\alpha,\beta) \;=\; \alpha_i \quad \forall i \in \mathcal{I}\,, \quad \text{and} \quad \sum_i \Psi_{ij}(\alpha,\beta) \;=\; \beta_j \quad \forall j \in \mathcal{J}\,, \tag{2.9}$$

with $\Psi_{ij}(\alpha,\beta) = 0$ for $i \nsim j$.

We quote a result from [21], which is used later. The proof of Lemma 3 in [21] assumes that the limits in (2.1) exist, in particular, $\sqrt{n}(N^n_j - n\nu_j) \to 0$ as $n \to \infty$. Nevertheless, the proof goes through under the weaker assumption that $N^n_j - n\nu_j = \mathfrak{o}(n)$.

Lemma 2.1 (Lemma 3 in [21]). *There exists a constant $M_0 > 0$ such that, the collection of sets $\breve{\mathfrak{X}}^n$ defined by*

$$\breve{\mathfrak{X}}^n \;:=\; \bigl\{x \in \mathbb{Z}^I_+ \,:\, \|x - nx^*\| \le M_0\, n\bigr\}\,, \tag{2.10}$$

satisfies $\breve{\mathfrak{X}}^n \subset \mathfrak{X}^n$ for all $n \in \mathbb{N}$. Moreover,

$$\Psi(x-q, N^n - y) \in \mathbb{Z}^{I\times J}_+ \qquad \forall\, q, y \in \Theta^n(x)\,, \quad \forall\, x \in \breve{\mathfrak{X}}^n\,.$$

Remark 2.1. Lemma 2.1 implies that if $x \in \breve{\mathfrak{X}}^n$, then for any $q \in \mathbb{Z}^I_+$ and $y \in \mathbb{Z}^J_+$ satisfying $e\cdot q \wedge e\cdot y = 0$ and $e\cdot(x-q) = e\cdot(N^n - y) \ge 0$, we have $\Psi(x-q, N^n-y) \in \mathcal{Z}^n(x)$.

We need the following definition.

Definition 2.2. We fix some open ball $\breve{B}$ centered at the origin, such that $n(\breve{B}+x^*) \subset \breve{\mathfrak{X}}^n$ for all $n \in \mathbb{N}$. The *jointly work conserving action set $\breve{\mathcal{Z}}^n(x)$ at x* is defined as the subset of $\mathcal{Z}^n(x)$, which satisfies

$$\breve{\mathcal{Z}}^n(x) \;:=\; \begin{cases} \bigl\{z \in \mathcal{Z}^n(x) : \;\; e\cdot q(x,z) \wedge e\cdot y^n(z) \;=\; 0\bigr\} & \text{if } \; x \in n(\breve{B}+x^*)\,,\\ \mathcal{Z}^n(x) & \text{otherwise,}\end{cases}$$

with q and y^n as in (2.5). We also define the associated admissible policies by

$$\breve{\mathfrak{Z}}^n \;:=\; \bigl\{Z^n \in \mathfrak{Z}^n \,:\, Z^n(t) \in \breve{\mathcal{Z}}^n\bigl(X^n(t)\bigr) \quad \forall\, t \ge 0\bigr\}\,,$$

$$\boldsymbol{\mathfrak{Z}} \;:=\; \{Z^n \in \breve{\mathfrak{Z}}^n\,,\; n \in \mathbb{N}\}\,.$$

We refer to the policies in $\boldsymbol{\mathfrak{Z}}$ as *eventually jointly work conserving* (EJWC).

The ball $\breve{B}$ is fixed in Definition 2.2 only for convenience. We could instead adopt a more general definition of $\boldsymbol{\mathfrak{Z}}$, as explained in Remark 2.1 in [15]. The EJWC condition plays a crucial role in the derivation of the controlled diffusion limit. Therefore, the convergence of mean empirical measures of the controlled diffusion-scaled state process, and thus, also the lower and upper bounds for asymptotic optimality are established for sequences $\{Z^n,\, n \in \mathbb{N}\} \subset \boldsymbol{\mathfrak{Z}}$.

2.3. The diffusion-scaled processes

Let x^* and z^* be as in (2.3). We define the diffusion-scaled processes $\hat{Z}^n$, $\hat{X}^n$, $\hat{Q}^n$, and $\hat{Y}^n$, by

$$\begin{aligned} \hat{X}^n_i(t) \;&:=\; \frac{1}{\sqrt{n}}(X^n_i(t) - nx^*_i)\,, & \hat{Z}^n_{ij}(t) \;&:=\; \frac{1}{\sqrt{n}}(Z^n_{ij}(t) - nz^*_{ij})\,,\\ \hat{Q}^n_i(t) \;&:=\; \frac{1}{\sqrt{n}}Q^n_i(t)\,, & \hat{Y}^n_j(t) \;&:=\; \frac{1}{\sqrt{n}}Y^n_j(t)\,. \end{aligned} \tag{2.11}$$

Let

$$\hat{M}^n_{A,i}(t) \;:=\; \frac{1}{\sqrt{n}}(A^n_i(\lambda^n_i t) - \lambda^n_i t),$$

$$\hat{M}^n_{S,ij}(t) \;:=\; \frac{1}{\sqrt{n}}\left(S^n_{ij}\left(\mu^n_{ij}\int_0^t Z^n_{ij}(s)\mathrm{d}s\right) - \mu^n_{ij}\int_0^t Z^n_{ij}(s)\mathrm{d}s\right),$$

$$\hat{M}^n_{R,i}(t) \;:=\; \frac{1}{\sqrt{n}}\left(R^n_i\left(\gamma^n_i\int_0^t Q^n_i(s)\mathrm{d}s\right) - \gamma^n_i\int_0^t Q^n_i(s)\mathrm{d}s\right).$$

These are square integrable martingales w.r.t. the filtration $\boldsymbol{\mathcal{F}}^n$, with quadratic variations

$$\langle \hat{M}^n_{A,i}\rangle(t) \;:=\; \frac{\lambda^n_i}{n}t\,, \quad \langle \hat{M}^n_{S,ij}\rangle(t) \;:=\; \frac{\mu^n_{ij}}{n}\int_0^t Z^n_{ij}(s)\mathrm{d}s\,, \quad \langle \hat{M}^n_{R,i}\rangle(t) \;:=\; \frac{\gamma^n_i}{n}\int_0^t Q^n_i(s)\mathrm{d}s\,.$$

Let $\widehat{M}^n(t) := \hat{M}^n_{A,i}(t) - \sum_{j\in\mathcal{J}(i)} \hat{M}^n_{S,ij}(t) - \hat{M}^n_{R,i}(t)$. By (2.6), we can write $\hat{X}^n_i(t)$ as

$$\hat{X}^n_i(t) \;=\; \hat{X}^n_i(0) + \ell^n_i t - \sum_{j\in\mathcal{J}(i)} \mu^n_{ij}\int_0^t \hat{Z}^n_{ij}(s)\mathrm{d}s - \gamma^n_i\int_0^t \hat{Q}^n_i(s)\mathrm{d}s + \widehat{M}^n(t)\,, \tag{2.12}$$

where $\ell^n = (\ell^n_1,\dots,\ell^n_I)^{\mathsf{T}}$ is defined as

$$\ell^n_i \;:=\; \frac{1}{\sqrt{n}}\biggl(\lambda^n_i - \sum_{j\in\mathcal{J}(i)} \mu^n_{ij} z^*_{ij} n\biggr)\,.$$

Under the assumptions on the parameters in (2.1) and the first constraint in the LP, it holds that

$$\ell^n_i \xrightarrow[n\to\infty]{} \ell_i \;:=\; \hat{\lambda}_i - \sum_{j\in\mathcal{J}(i)} \hat{\mu}_{ij} z^*_{ij}\,.$$

We let $\ell := (\ell_1,\dots,\ell_I)^{\mathsf{T}}$.

By (2.3), (2.4), and (2.11), we obtain the balance equations

$$\begin{aligned} \hat{X}^n_i(t) \;&=\; \hat{Q}^n_i(t) + \sum_{j\in\mathcal{J}(i)} \hat{Z}^n_{ij}(t) \qquad \forall\, i\in\mathcal{I}\,,\\ \hat{Y}^n_j(t) + \sum_{i\in\mathcal{I}(j)} \hat{Z}^n_{ij}(t) \;&=\; 0 \qquad \forall\, j\in\mathcal{J}\,. \end{aligned} \tag{2.13}$$

Definition 2.3. For each $x\in\mathbb{Z}^I_+$ and $z\in\mathcal{Z}^n(x)$, we define

$$\begin{aligned} \tilde{x}^n \;&=\; \tilde{x}^n(x) \;:=\; x - nx^*\,, \qquad \hat{x}^n \;=\; \hat{x}^n(x) \;:=\; \frac{\tilde{x}^n(x)}{\sqrt{n}}\,, \qquad \hat{z}^n(z) \;:=\; \frac{z-nz^*}{\sqrt{n}}\,,\\ \hat{q}^n(x,z) \;&:=\; \frac{q(x,z)}{\sqrt{n}}\,, \qquad \hat{y}^n(z) \;:=\; \frac{y^n(z)}{\sqrt{n}}\,, \qquad \hat{\vartheta}^n(x,z) \;:=\; e\cdot\hat{q}^n(x,z)\wedge e\cdot\hat{y}^n(z)\,, \end{aligned} \tag{2.14}$$

with $q(x,z)$, $y^n(z)$ as in (2.5). We also let

$$\mathcal{S}^n \;:=\; \bigl\{\hat{x}^n(x) :\, x\in\mathbb{Z}^I_+\bigr\}\,, \qquad \breve{\mathcal{S}}^n \;:=\; \bigl\{\hat{x}^n(x) :\, x\in\breve{\mathcal{X}}^n\bigr\}$$

and

$$\hat{\mathcal{Z}}^n(\hat{x}) \;:=\; \{\hat{z}^n(z) :\, z\in\mathcal{Z}^n(\sqrt{n}\hat{x}+nx^*)\} \qquad \hat{x}\in\mathcal{S}^n\,.$$

Abusing the notation, we also write

$$\hat{q}^n_i(\hat{x},\hat{z}) \;=\; \hat{q}^n_i(\hat{x}^n,\hat{z}^n) \;=\; \hat{x}^n_i - \sum_{j\in\mathcal{J}(i)} \hat{z}^n_{ij} \quad \text{for} \quad i\in\mathcal{I}\,, \tag{2.15}$$

and

$$\hat{y}^n_j(\hat{z}) \;=\; \hat{y}^n_j(\hat{z}^n) \;=\; \frac{N^n_j - n\sum_{i\in\mathcal{I}(j)} z^*_{ij}}{\sqrt{n}} - \sum_{i\in\mathcal{I}(j)} \hat{z}^n_{ij} \quad \text{for} \quad j\in\mathcal{J}\,. \tag{2.16}$$

Lemma 2.2. *There exists a constant $\tilde{M}_0 > 0$ such that for any $z \in \breve{\mathcal{Z}}^n(x)$, $x \in \mathbb{Z}^I_+$, and $n \in \mathbb{N}$, we have*

$$\max\left\{\max_{(i,j)\in\mathcal{E}} |\hat{z}^n_{ij}(z)|,\ \|\hat{q}^n(x,z)\|,\ \|\hat{y}^n(z)\|,\ \hat{\vartheta}^n(x,z)\right\} \;\le\; \tilde{M}_0\,\|\hat{x}^n(x)\|\,.$$

Proof. Note that

$$\|\hat{q}^n(x,z)\| = \hat{\vartheta}^n(x,z) + \bigl(e\cdot\hat{x}^n(x)\bigr)^+\,,\quad\text{and}\quad \|\hat{y}^n(z)\| = \hat{\vartheta}^n(x,z) + \bigl(e\cdot\hat{x}^n(x)\bigr)^- \tag{2.17}$$

for all $x \in \mathbb{Z}^I_+$ and $z \in \mathcal{Z}^n(x)$. Therefore, there exist probability vectors $p^c \in [0,1]^I$ and $p^s \in [0,1]^J$ such that $\hat{q}^n = \bigl(\hat{\vartheta}^n + (e\cdot\hat{x}^n)^+\bigr)p^c$ and $\hat{y}^n = \bigl(\hat{\vartheta}^n + (e\cdot\hat{x}^n)^-\bigr)p^s$. By the linearity of the map Ψ and Lemma 2.1, it easily follows that

$$\hat{z}^n \;=\; \Psi(\hat{x}^n - \hat{q}^n, -\hat{y}^n) \;=\; \Psi\bigl(\hat{x}^n - (e\cdot\hat{x}^n)^+p^c, -(e\cdot\hat{x}^n)^-p^s\bigr) - \hat{\vartheta}^n\,\Psi(p^c,p^s)\,. \tag{2.18}$$

If $x \notin \breve{\mathfrak{X}}^n$, then $\|\hat{x}^n\| > M_0\sqrt{n}$ by (2.10). Since for some constant $C > 0$, it holds that $\|\hat{y}^n(z)\| \le C\sqrt{n}$ for all $n \in \mathbb{N}$, the same bound also holds for $\hat{\vartheta}^n(x,z)$. Thus if $x \notin \breve{\mathfrak{X}}^n$, we obtain the bound asserted in the lemma by (2.17) and (2.18).

On the other hand, if $x \in \breve{\mathfrak{X}}^n$ and $z \in \breve{\mathcal{Z}}^n(x)$, then $\hat{\vartheta}^n(x,z) = 0$, and again the assertion of the lemma follows by (2.17) and (2.18). This completes the proof. □

Definition 2.4. We define the operator $\mathcal{A}^n\colon \mathcal{C}^2(\mathbb{R}^I) \to \mathcal{C}(\mathbb{R}^I, \mathbb{R}^{I\times J})$ by

$$\mathcal{A}^n f(\hat{x},\hat{z}) \;\coloneqq\; \sum_{i\in\mathcal{I}}\Bigl(\mathcal{A}^n_{i,1}(\hat{x}_i,\hat{z})\,\partial_i f(\hat{x}) + \mathcal{A}^n_{i,2}(\hat{x}_i,\hat{z})\,\partial_{ii} f(\hat{x})\Bigr)\,,\qquad f \in \mathcal{C}^2(\mathbb{R}^I)\,,$$

where $\partial_i \coloneqq \frac{\partial}{\partial x_i}$ and $\partial_{ij} \coloneqq \frac{\partial^2}{\partial x_i \partial x_j}$, and

$$\mathcal{A}^n_{i,1}(\hat{x}_i,\hat{z}) \;\coloneqq\; \ell^n_i - \sum_{j\in\mathcal{J}(i)} \mu^n_{ij}\hat{z}_{ij} - \gamma^n_i\Bigl(\hat{x}_i - \sum_{j\in\mathcal{J}(i)} \hat{z}_{ij}\Bigr)\,,$$

$$\mathcal{A}^n_{i,2}(\hat{x}_i,\hat{z}) \;\coloneqq\; \frac{1}{2}\left[\frac{\lambda^n_i}{n} + \sum_{j\in\mathcal{J}(i)} \mu^n_{ij} z^*_{ij} + \frac{1}{\sqrt{n}}\sum_{j\in\mathcal{J}(i)} \mu^n_{ij}\hat{z}_{ij} + \frac{\gamma^n_i}{\sqrt{n}}\Bigl(\hat{x}_i - \sum_{j\in\mathcal{J}(i)} \hat{z}_{ij}\Bigr)\right].$$

By the Kunita–Watanabe formula for semi-martingales (see, e.g., [37, Theorem 26.7]), we have

$$f(\hat{X}^n(t)) \;=\; f(\hat{X}^n(0)) + \int_0^t \mathcal{A}^n f\bigl(\hat{X}^n(s), \hat{Z}^n(s)\bigr)\,\mathrm{d}s + \sum_{s\le t} \mathcal{D} f(\hat{X}^n, s) \qquad \forall\, f \in \mathcal{C}^2(\mathbb{R}^I)\,, \tag{2.19}$$

for any admissible diffusion-scaled policy $\hat{Z}^n$, where

$$\mathcal{D} f(\hat{X}^n, s) \;\coloneqq\; \Delta f(\hat{X}^n(s)) - \sum_{i\in\mathcal{I}} \partial_i f(\hat{X}^n(s-))\Delta\hat{X}^n_i(s) - \frac{1}{2}\sum_{i,i'\in\mathcal{I}} \partial_{ii'} f(\hat{X}^n(s-))\Delta\hat{X}^n_i(s)\Delta\hat{X}^n_{i'}(s)\,. \tag{2.20}$$

2.3.1. Control parameterization

Definition 2.5. Let $\mathfrak{X}^n \coloneqq \{(\hat{x},\hat{z})\colon \hat{x} \in \mathcal{S}^n\,,\ \hat{z} \in \hat{\mathcal{Z}}^n(\hat{x})\}$. For each $(\hat{x},\hat{z}) \in \mathfrak{X}^n$, we define

$$u^c_i(\hat{x},\hat{z}) \;=\; u^{c,n}_i(\hat{x},\hat{z}) \;\coloneqq\; \begin{cases} \frac{\hat{q}^n_i(\hat{x},\hat{z})}{e\cdot\hat{q}^n(\hat{x},\hat{z})} & \text{if } e\cdot\hat{q}^n(\hat{x},\hat{z}) > 0\,,\\ e_I & \text{otherwise,}\end{cases} \qquad i \in \mathcal{I}\,,\quad t \ge 0\,,$$

and

$$u^s_j(\hat z) \;=\; u^{s,n}_j(\hat z) \;\coloneqq\; \begin{cases} \frac{\hat y^n_j(\hat z)}{e\cdot \hat y^n(\hat z)} & \text{if } e\cdot\hat y^n(\hat z)>0\,,\\ e_J & \text{otherwise,}\end{cases} \qquad j\in\mathcal{J}\,,\quad t\ge 0\,.$$

Let $u(\hat x,\hat z)\coloneqq (u^c(\hat x,\hat z),u^s(\hat z))$. Then $u(\hat x,\hat z)$ belongs to the set

$$\mathbb{U} \;\coloneqq\; \bigl\{u=(u^c,u^s)\in\mathbb{R}^I_+\times\mathbb{R}^J_+ :\, e\cdot u^c = e\cdot u^s = 1\bigr\}\,. \tag{2.21}$$

We also define the processes

$$U^{c,n}_i(t) \;\coloneqq\; u^c_i\bigl(\hat X^n(t),\hat Z^n(t)\bigr)\,, \qquad U^{s,n}_i(t) \;\coloneqq\; u^s_i\bigl(\hat Z^n(t)\bigr)\,,$$

and $U^n\coloneqq (U^{c,n},U^{s,n})$, with $U^{c,n}\coloneqq (U^{c,n}_1,\dots,U^{c,n}_I)^{\mathsf{T}}$, and $U^{s,n}\coloneqq (U^{s,n}_1,\dots,U^{s,n}_J)^{\mathsf{T}}$.

The process $U^{c,n}_i(t)$ represents the proportion of the total queue length in the network at queue i at time t, while $U^{s,n}_j(t)$ represents the proportion of the total idle servers in the network at station j at time t. Given $Z^n\in\mathfrak{Z}^n$ the process U^n is uniquely determined and lives in the set $\mathbb{U}$.

For $u\in\mathbb{U}$, let $\widehat\Psi[u]\colon\mathbb{R}^I\to\mathbb{R}^{\mathcal{G}}$ be defined by

$$\widehat\Psi[u](x) \;\coloneqq\; \Psi(x-(e\cdot x)^+u^c, -(e\cdot x)^-u^s)\,, \tag{2.22}$$

where Ψ is as in (2.9).

We define the operator $\breve{\mathcal{A}}^n\colon\mathcal{C}^2(\mathbb{R}^I)\to\mathcal{C}(\mathbb{R}^I,\mathbb{U})$ by

$$\breve{\mathcal{A}}^n f(\hat x,u) \;\coloneqq\; \sum_{i\in\mathcal{I}}\Bigl(\breve{\mathcal{A}}^n_{i,1}(\hat x,u)\,\partial_i f(\hat x) + \breve{\mathcal{A}}^n_{i,2}(\hat x,u)\,\partial_{ii} f(\hat x)\Bigr)\,,$$

where

$$\breve{\mathcal{A}}^n_{i,1}(\hat x,u) \;\coloneqq\; \ell^n_i - \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}\widehat\Psi_{ij}[u](\hat x) - \gamma^n_i(e\cdot\hat x)^+u^c_i\,,$$

$$\breve{\mathcal{A}}^n_{i,2}(\hat x,u) \;\coloneqq\; \frac{1}{2}\Biggl(\frac{\lambda^n_i}{n} + \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}z^*_{ij} + \frac{1}{\sqrt n}\sum_{j\in\mathcal{J}(i)}\mu^n_{ij}\widehat\Psi_{ij}[u](\hat x) + \frac{\gamma^n_i}{\sqrt n}\bigl((e\cdot\hat x)^+u^c_i\bigr)\Biggr)\,.$$

The following lemma is a result of a simple calculation based on the definitions above. Recall the definitions of $\breve B$, $\breve{\mathcal{Z}}^n$, and $\breve{\mathbb{S}}^n$ from Definitions 2.2 and 2.3.

Lemma 2.3. *Let $u=u(\hat x,\hat z)\colon\mathfrak{X}^n\to\mathbb{U}$ denote the map given in Definition 2.5. Then for $f\in\mathcal{C}^2_c(\sqrt n\breve B)$, we have*

$$\breve{\mathcal{A}}^n f\bigl(\hat x,u(\hat x,\hat z)\bigr) = \mathcal{A}^n f\bigl(\hat x,\hat z\bigr)\,, \qquad \forall\,\hat x\in\breve{\mathbb{S}}^n\cap\sqrt n\breve B\,,\ \forall\,\hat z\in\breve{\mathcal{Z}}^n(\sqrt n\hat x+nx^*)\,.$$

2.4. The diffusion limit

Consider the I-dimensional controlled diffusion given by the Itô equation

$$dX_t \;=\; b(X_t,U_t)\,\mathrm{d}t + \Sigma\,\mathrm{d}W_t\,, \tag{2.24}$$

where W is an I-dimensional standard Wiener process. The drift $b\colon\mathbb{R}^I\times\mathbb{U}\to\mathbb{R}^I$ takes the form

$$b_i(x,u) \;=\; b_i(x,(u^c,u^s)) \;\coloneqq\; \ell_i - \sum_{j\in\mathcal{J}(i)}\mu_{ij}\widehat\Psi_{ij}[u](x) - \gamma_i(e\cdot x)^+u^c_i \qquad \forall\,i\in\mathcal{I}\,,$$

where $\widehat{\Psi}_{ij}[u]$ is as in (2.22). Also $\Sigma := \operatorname{diag}\bigl(\sqrt{2\lambda_1}, \dots, \sqrt{2\lambda_I}\bigr)$.

The control process U takes values in $\mathbb{U}$, defined in (2.21), and $U_t(\omega)$ is jointly measurable in $(t, \omega) \in [0, \infty) \times \Omega$. Moreover, it is *non-anticipative*, i.e., for $s < t$, $W_t - W_s$ is independent of

$$\mathfrak{F}_s := \text{the completion of } \sigma\{X_0, U_r, W_r,\ r \le s\} \text{ relative to } (\mathfrak{F}, \mathbb{P})\,.$$

Such a process U is called an *admissible control*. Let $\mathfrak{U}$ denote the set of all admissible controls. Recall that a control is called *Markov* if $U_t = v(t, X_t)$ for a measurable map $v\colon \mathbb{R}_+ \times \mathbb{R}^I \to \mathbb{U}$, and it is called *stationary Markov* if v does not depend on t, i.e., $v\colon \mathbb{R}^I \to \mathbb{U}$. Let $\mathfrak{U}_{\mathrm{SM}}$ denote the set of stationary Markov controls. Recall also that a control $v \in \mathfrak{U}_{\mathrm{SM}}$ is called *stable* if the controlled process is positive recurrent. We denote the set of such controls by $\mathfrak{U}_{\mathrm{SSM}}$. Let

$$\mathcal{L}^u f(x) := \sum_{i\in\mathcal{I}} \bigl[\lambda_i\, \partial_{ii} f(x) + b_i(x,u)\, \partial_i f(x)\bigr]\,, \quad u \in \mathbb{U}\,, \tag{2.25}$$

denote the extended controlled generator of the diffusion in (2.24).

In [15], a leaf elimination algorithm was developed to obtain an explicit expression for the drift $b(x,u)$. This plays an important role in understanding the recurrence properties of the controlled diffusion. See also Remark 4.2 and Example 4.4 in [15]. We quote this result as follows.

Lemma 2.4 (Lemma 4.3 in [15]). *The drift $b(x,u) = b(x,(u^c,u^s))$ in the limiting diffusion X in (2.24) can be expressed as*

$$b(x,u) = \ell - B_1(x - (e\cdot x)^+ u^c) + (e\cdot x)^- B_2 u^s - (e\cdot x)^+ \varGamma u^c\,, \tag{2.26}$$

where B_1 is a lower-diagonal $I\times I$ matrix with positive diagonal elements, B_2 is an $I\times J$ matrix and $\varGamma = \operatorname{diag}\{\gamma_1, \dots, \gamma_I\}$.

The drift in (2.26) takes the form

$$b_i\bigl(x,u\bigr) = \ell_i - \mu_{ij_i} x_i + \tilde{b}_i(x_1, \dots, x_{i-1}) + \tilde{F}_i\bigl((e\cdot x)^+ u^c, (e\cdot x)^- u^s\bigr) - \gamma_i\, (e\cdot x)^+ u_i^c\,, \tag{2.27}$$

where $j_i \in \mathcal{J}$, $i \sim j_i$, is the unique server-pool node corresponding to i when customer node i is removed by the leaf elimination algorithm (see Section 4.1 in [15]). Two things are important to note: (a) $\tilde{F}_i$ is a linear function, and (b) $\mu_{ij_i} > 0$ (since $i \sim j_i$).

Under EJWC policies, convergence in distribution of the diffusion-scaled processes $\hat{X}^n$ to the limiting diffusion X in (2.24) follows by [21, Proposition 3] for certain classes of networks. The fact that (2.24) can be viewed as a limit of the diffusion-scaled process $\hat{X}^n$ is also indicated by the following lemma.

Lemma 2.5. *We have*

$$\breve{\mathcal{A}}^n_{i,1} \xrightarrow[n\to\infty]{} b_i\,, \quad \textit{and} \quad \breve{\mathcal{A}}^n_{i,2} \xrightarrow[n\to\infty]{} \lambda_i$$

for $i \in \mathcal{I}$, uniformly over compact sets of $\mathbb{R}^I \times \mathbb{U}$. In particular, for any $f \in \mathcal{C}^2_c(\mathbb{R}^I)$ it holds that

$$\breve{\mathcal{A}}^n f(x,u) \xrightarrow[n\to\infty]{} \mathcal{L}^u f(x)\,.$$

Nevertheless, for the time being, we consider solutions of (2.24) as the formal limit of (2.12). Precise links of the n^{th} system model and (2.24) are established in Section 6.

Definition 2.6. Let $\|x\|_\beta \coloneqq (\beta_1 |x_1|^2 + \cdots + \beta_I |x_I|^2)^{1/2}$, with $\beta = (\beta_1, \dots, \beta_I)$ a positive vector. Throughout the paper, $\mathcal{V}_{\kappa,\beta}$, $\kappa \ge 1$, stands for a $\mathcal{C}^2(\mathbb{R}^I)$ function which agrees with $\|x\|_\beta^\kappa$ on the complement of the unit ball B in $\mathbb{R}^I$, i.e., $\mathcal{V}_{\kappa,\beta}(x) = \|x\|_\beta^\kappa$, for $x \in B^c$. Also, $\widetilde{\mathcal{V}}_{\epsilon,\beta}$, $\epsilon > 0$, is defined by

$$\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \coloneqq \exp\Bigl(\epsilon\|x\|_\beta^2\bigl(1 + \|x\|_\beta^2\bigr)^{-1/2}\Bigr)\,, \qquad x \in \mathbb{R}^I\,.$$

In addition, for $\delta > 0$, we define

$$\mathcal{K}_\delta \coloneqq \bigl\{x \in \mathbb{R}^I : |e \cdot x| > \delta|x|\bigr\}\,.$$

As shown in Theorem 4.1 of [15], the drift b in (2.27) has the following important structural property. For any $\kappa \ge 1$, there exists a function $\mathcal{V}_{\kappa,\beta}$ as in Definition 2.6, and positive constants c_i, $i = 0, 1, 2$, such that

$$b(x,u) \cdot \nabla\mathcal{V}_{\kappa,\beta}(x) \le c_0 - c_1\mathcal{V}_{\kappa,\beta}(x)\,\mathbb{1}_{\mathcal{K}_\delta^c}(x) + c_2\mathcal{V}_{\kappa,\beta}(x)\,\mathbb{1}_{\mathcal{K}_\delta}(x) \qquad \forall\,(x,u) \in \mathbb{R}^I \times \mathbb{U}\,. \tag{2.28}$$

Since the diffusion matrix is constant, it is evident that a similar estimate holds for $\mathcal{L}^u\mathcal{V}_{\kappa,\beta}$ uniformly over $u \in \mathbb{U}$. By a straightforward application of Itô's formula, this implies that for any $\kappa \ge 1$ there exists a constant C depending only on κ such that (see [15, Lemma 3.1 (c)])

$$\mathbb{E}_x^U\biggl[\int_0^T |X_s|^\kappa\,\mathrm{d}s\biggr] \le C\,|x|^\kappa + C\,\mathbb{E}_x^U\biggl[\int_0^T \bigl(1 + |e \cdot X_s|\bigr)^\kappa\,\mathrm{d}s\biggr] \qquad \forall\,T > 0\,, \quad \forall\,U \in \mathfrak{U}\,. \tag{2.29}$$

Moreover, it is shown in [15, Theorem 4.2] that there exists a stationary Markov control $\bar{v} \in \mathfrak{U}_{\mathrm{SM}}$ satisfying

$$\mathcal{L}^{\bar{v}}\mathcal{V}_{\kappa,\beta}(x) \le \bar{c}_0 - \bar{c}_1\mathcal{V}_{\kappa,\beta}(x) \qquad \forall x \in \mathbb{R}^I\,, \tag{2.30}$$

for any $\kappa \ge 1$, and positive constants $\bar{c}_0$ and $\bar{c}_1$ depending only on κ. As a consequence of (2.30), the diffusion under the control $\bar{v}$ is exponentially ergodic. A slight modification of that proof leads to the following theorem.

Theorem 2.1. *Provided that $\gamma_i > 0$ for some $i \in \mathcal{I}$, there exist $\epsilon > 0$, a positive vector $\beta \in \mathbb{R}^I$, and a stationary Markov control $\bar{v} \in \mathfrak{U}_{\mathrm{SM}}$ satisfying*

$$\mathcal{L}^{\bar{v}}\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \le \tilde{c}_0 - \tilde{c}_1\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \qquad \forall x \in \mathbb{R}^I\,, \tag{2.31}$$

for some positive constants $\tilde{c}_0$ and $\tilde{c}_1$.

The properties in (2.28), (2.30), and (2.31) are instrumental in showing that the optimal control problems defined in this paper are well posed.

3. Ergodic Control Problems

In this section, we consider two control objectives, which address the queueing (delay) and/or idleness costs in the system: (i) *unconstrained problem*, minimizing the queueing and idleness cost and (ii) *constrained problem*, minimizing the queueing cost while imposing a constraint on idleness. We state both problems for the n^{th} system and the limiting diffusion.

3.1. Ergodic control problems for the n^{th} system

The running cost is a function of the diffusion-scaled processes, which are related to the unscaled ones by (2.11). For simplicity, in all three cost minimization problems, we assume that the initial condition $X^n(0)$ is deterministic and $\hat{X}^n(0) \to x \in \mathbb{R}^I$ as $n \to \infty$. Let the running cost $\hat{r} \colon \mathbb{R}^I_+ \times \mathbb{R}^J_+ \to \mathbb{R}_+$ be defined by

$$\hat{r}(\hat{q}, \hat{y}) \;=\; \sum_{i \in \mathcal{I}} \xi_i \hat{q}_i^m + \sum_{j \in \mathcal{J}} \zeta_j \hat{y}_j^m\,, \quad \hat{q} \in \mathbb{R}^I_+\,,\ \hat{y} \in \mathbb{R}^J_+\,, \quad \text{for some } m \ge 1\,, \tag{3.1}$$

where $\xi = (\xi_1, \dots, \xi_I)^{\mathsf{T}}$ is a positive vector and $\zeta = (\zeta_1, \dots, \zeta_J)^{\mathsf{T}}$ is a nonnegative vector. In the case $\zeta \equiv 0$, only the queueing cost is minimized. We denote by $\mathbb{E}^{Z^n}$ the expectation operator under an admissible policy Z^n.

(P1) (*unconstrained problem*) The running cost penalizes the queueing and idleness. Let $\hat{r}(q, y)$ be the running cost function as defined in (3.1). Here $\zeta > 0$. Given an initial state $X^n(0)$, and an admissible scheduling policy $Z^n \in \breve{\mathfrak{Z}}^n$, we define the diffusion-scaled cost criterion by

$$J\big(\hat{X}^n(0), Z^n\big) \;\coloneqq\; \limsup_{T \to \infty} \frac{1}{T}\, \mathbb{E}^{Z^n} \left[\int_0^T \hat{r}\big(\hat{Q}^n(s), \hat{Y}^n(s)\big)\, \mathrm{d}s \right]. \tag{3.2}$$

The associated cost minimization problem becomes

$$\hat{V}^n(\hat{X}^n(0)) \;\coloneqq\; \inf_{Z^n \in \breve{\mathfrak{Z}}^n} J\big(\hat{X}^n(0), Z^n\big)\,.$$

(P2) (*constrained problem*) The objective here is to minimize the queueing cost while imposing idleness constraints on the server pools. Let $\hat{r}_{\mathsf{o}}(q)$ be the running cost function corresponding to $\hat{r}$ in (3.1) with $\zeta \equiv 0$. The diffusion-scaled cost criterion $J_{\mathsf{o}}\big(\hat{X}^n(0), Z^n\big)$ is defined analogously to (3.2) with running cost $\hat{r}_{\mathsf{o}}(\hat{Q}^n(s))$, that is,

$$J_{\mathsf{o}}\big(\hat{X}^n(0), Z^n\big) \;\coloneqq\; \limsup_{T \to \infty} \frac{1}{T}\, \mathbb{E}^{Z^n} \left[\int_0^T \hat{r}_{\mathsf{o}}\big(\hat{Q}^n(s)\big)\, \mathrm{d}s \right].$$

Also define

$$J_{\mathsf{c},j}\big(\hat{X}^n(0), Z^n\big) \;\coloneqq\; \limsup_{T \to \infty} \frac{1}{T}\, \mathbb{E}^{Z^n} \left[\int_0^T \big(\hat{Y}^n_j(s)\big)^{\tilde{m}}\, \mathrm{d}s \right], \qquad j \in \mathcal{J}\,,$$

with $\tilde{m} \ge 1$. The associated cost minimization problem becomes

$$\begin{aligned} \hat{V}^n_{\mathsf{c}}(\hat{X}^n(0)) \;&\coloneqq\; \inf_{Z^n \in \breve{\mathfrak{Z}}^n} J_{\mathsf{o}}\big(\hat{X}^n(0), Z^n\big)\,, \\ &\text{subject to} \quad J_{\mathsf{c},j}\big(\hat{X}^n(0), Z^n\big) \;\le\; \updelta_j\,, \quad j \in \mathcal{J}\,, \end{aligned} \tag{3.3}$$

where $\updelta = (\updelta_1, \dots, \updelta_J)^{\mathsf{T}}$ is a positive vector.

We refer to $\hat{V}^n(\hat{X}^n(0))$ and $\hat{V}^n_{\mathsf{c}}(\hat{X}^n(0))$ as the diffusion-scaled optimal values for the n^{th} system given the initial state $X^n(0)$, for (P1) and (P2), respectively.

Remark 3.1. We choose running costs of the form (3.1) mainly to simplify the exposition. However, all the results of this paper still hold for more general classes of functions. Let $h_{\mathsf{o}} \colon \mathbb{R}^I \to \mathbb{R}_+$ be a convex function satisfying $h_{\mathsf{o}}(x) \ge c_1 |x|^m + c_2$ for some $m \ge 1$ and constants $c_1 > 0$ and $c_2 \in \mathbb{R}$, and $h \colon \mathbb{R}^I \to \mathbb{R}_+$, $h_i \colon \mathbb{R} \to \mathbb{R}_+$, $i \in \mathcal{I}$, be convex functions that have at most polynomial growth. Then we can choose $\hat{r}(q, y) = h_{\mathsf{o}}(q) + h(y)$ for the unconstrained problem, and $h_i(y_i)$ as the functions in the constraints in (3.3) (with $\hat{r}_{\mathsf{o}} = h_{\mathsf{o}}$). Analogous running costs can of course be used in the corresponding control problems for the limiting diffusion, which are presented later in Section 3.2.

3.2. Ergodic control problems for the limiting diffusion

We state the two problems which correspond to (P1)–(P2) in Section 3.1 for the controlled diffusion in (2.24). Let $r\colon \mathbb{R}^I \times \mathbb{U} \to \mathbb{R}$ be defined by

$$r(x,u) \;=\; r\bigl(x,(u^c,u^s)\bigr) \;\coloneqq\; \hat{r}\bigl((e\cdot x)^+ u^c, (e\cdot x)^- u^s\bigr)\,,$$

with $\hat{r}$ as in (3.1), that is,

$$r(x,u) \;=\; [(e\cdot x)^+]^m \sum_{i\in\mathcal{I}} \xi_i (u_i^c)^m + [(e\cdot x)^-]^m \sum_{j\in\mathcal{J}} \zeta_j (u_j^s)^m, \quad m\ge 1\,, \tag{3.4}$$

for the given $\xi = (\xi_1,\dots,\xi_I)^{\mathsf{T}}$ and $\zeta = (\zeta_1,\dots,\zeta_J)^{\mathsf{T}}$ in (3.1). Let the ergodic cost associated with the controlled diffusion X and the running cost r be defined as

$$J_{x,U}[r] \;\coloneqq\; \limsup_{T\to\infty} \frac{1}{T}\, \mathbb{E}_x^U \left[\int_0^T r(X_t,U_t)\,\mathrm{d}t\right], \quad U\in\mathfrak{U}\,.$$

(P1′) (*unconstrained problem*) The running cost function $r(x,u)$ is as in (3.4) with $\zeta > 0$. The ergodic control problem is then defined as

$$\varrho^*(x) \;\coloneqq\; \inf_{U\in\mathfrak{U}} J_{x,U}[r]\,.$$

(P2′) (*constrained problem*) The running cost function $r_{\mathsf{o}}(x,u)$ is as in (3.4) with $\zeta\equiv 0$. Also define

$$r_j(x,u) \;\coloneqq\; [(e\cdot x)^- u_j^s]^{\tilde{m}}\,, \quad j\in\mathcal{J}\,, \tag{3.5}$$

with $\tilde{m}\ge 1$, and let $\delta = (\delta_1,\dots,\delta_J)$ be a positive vector. The ergodic control problem under idleness constraints is defined as

$$\varrho_{\mathsf{c}}^*(x) \;\coloneqq\; \inf_{U\in\mathfrak{U}} J_{x,U}[r_{\mathsf{o}}]\,, \qquad \text{subject to} \quad J_{x,U}[r_j] \;\le\; \delta_j\,, \quad j\in\mathcal{J}\,.$$

The quantities $\varrho^*(x)$ and $\varrho_{\mathsf{c}}^*(x)$ are called the optimal values of the ergodic control problems (P1′) and (P2′), respectively, for the controlled diffusion process X with initial state x. Note that as is shown in Section 3 of [14] and Sections 3 and 5.4 of [15], the optimal values $\varrho^*(x)$ and $\varrho_{\mathsf{c}}^*(x)$ do not depend on $x\in\mathbb{R}^I$, and thus we remove this dependence in the results stated in Section 3.3.

Let $\mathcal{G}$ denote the set of ergodic occupation measures corresponding to controls in $\mathfrak{U}_{\mathrm{SSM}}$, that is,

$$\mathcal{G} \;\coloneqq\; \left\{\pi\in\mathcal{P}(\mathbb{R}^I\times\mathbb{U}) : \int_{\mathbb{R}^I\times\mathbb{U}} \mathcal{L}^u f(x)\,\pi(\mathrm{d}x,\mathrm{d}u) = 0 \quad \forall\, f\in\mathcal{C}_c^\infty(\mathbb{R}^I)\right\},$$

where $\mathcal{L}^u f(x)$ is the controlled extended generator of the diffusion X given in (2.25). The restriction of the ergodic control problem with running cost r to stable stationary Markov controls is equivalent to minimizing $\pi(r) = \int_{\mathbb{R}^I\times\mathbb{U}} r(x,u)\,\pi(\mathrm{d}x,\mathrm{d}u)$ over all $\pi\in\mathcal{G}$. If the infimum is attained in $\mathcal{G}$, then we say that the ergodic control problem is *well posed*, and we refer to any $\bar{\pi}\in\mathcal{G}$ that attains this infimum as an *optimal ergodic occupation measure.*

The characterization of the optimal solutions to the ergodic control problems (P1′)–(P2′) has been thoroughly studied in [14] and [15]. We refer the reader to these papers for relevant results used in the proof of asymptotic optimality which follows in the next section.

3.3. Asymptotic optimality results

We summarize here the main results on asymptotic optimality, which assert that the values of the two ergodic control problems in the diffusion scale converge to the values of the corresponding ergodic control problems for the limiting diffusion, respectively. The proofs of the asymptotic optimality are given in Section 7.

Recall the definitions of J, J_{o}, $\hat{V}^n$, and $\hat{V}^n_{\mathsf{c}}$ in (P1)–(P2), and the definitions of ϱ^* and ϱ^*_{c} in (P1′)–(P2′).

Theorem 3.1 (unconstrained problem). *Suppose that $\gamma_i > 0$ for some $i \in \mathcal{I}$. Then the following are true.*

(i) (*lower bound*) *For any sequence $\{Z^n,\ n \in \mathbb{N}\} \subset \mathfrak{Z}$, the diffusion-scaled cost in (3.2) satisfies*

$$\liminf_{n\to\infty}\ J\bigl(\hat{X}^n(0), \hat{Z}^n\bigr) \ \ge\ \varrho^*\,.$$

(ii) (*upper bound*) $\limsup\limits_{n\to\infty}\ \hat{V}^n(\hat{X}^n(0)) \ \le\ \varrho^*\,.$

Theorem 3.2 (constrained problem). *Under the assumptions of Theorem 3.1, we have the following:*

(i) (*lower bound*) *Suppose that under a sequence $\{Z^n,\ n \in \mathbb{N}\} \subset \mathfrak{Z}$, the constraint in (3.3) is satisfied for all sufficiently large $n \in \mathbb{N}$. Then*

$$\liminf_{n\to\infty}\ J_{\mathsf{o}}\bigl(\hat{X}^n(0), \hat{Z}^n\bigr) \ \ge\ \varrho^*_{\mathsf{c}}\,,$$

and as a result we have that $\liminf\limits_{n\to\infty}\ \hat{V}^n_{\mathsf{c}}(\hat{X}^n(0)) \ge \varrho^*_{\mathsf{c}}\,.$

(ii) (*upper bound*) *For any $\epsilon > 0$, there exists a sequence $\{Z^n,\ n \in \mathbb{N}\} \subset \mathfrak{Z}$ such that the constraint in (3.3) is feasible for all sufficiently large n, and*

$$\limsup_{n\to\infty}\ J_{\mathsf{o}}\bigl(\hat{X}^n(0), \hat{Z}^n\bigr) \ \le\ \varrho^*_{\mathsf{c}} + \epsilon\,.$$

Consequently, we have that $\limsup\limits_{n\to\infty}\ \hat{V}^n_{\mathsf{c}}(\hat{X}^n(0)) \ \le\ \varrho^*_{\mathsf{c}}\,.$

4. BQBS Stability and Fairness

4.1. BQBS stable networks

It follows by (2.29) that the controlled diffusion limit for multiclass multi-pool networks have the following property. If under some admissible control (admissible scheduling policy) the mean empirical value of some power $\kappa \ge 1$ of the queueing and idleness processes is bounded, then the corresponding mean empirical value of the state process also remains bounded. This property also holds for the diffusion-scaled processes in the n^{th} system, as shown later in Proposition 6.1.

There is however a large class of networks that share a more specific property, namely that the average value of any moment of a state process, is controlled by the average value of the corresponding moment of the queueing process alone. More precisely, the limiting diffusion of this class of networks satisfies

$$\mathbb{E}^U_x\biggl[\int_0^T |X_s|^\kappa\,\mathrm{d}s\biggr] \ \le\ C\,|x|^\kappa + C\,\mathbb{E}^U_x\biggl[\int_0^T \bigl[1 + \bigl(e\cdot X_s\bigr)^+\bigr]^\kappa\,\mathrm{d}s\biggr] \qquad \forall\,T>0\,,\quad \forall\,U\in\mathfrak{U}\,, \quad (4.1)$$

for any $\kappa \ge 1$, and for a constant C which depends only on κ. We refer to the class of networks which satisfy (4.1) as *bounded-queue, bounded-state* (BQBS) stable.

Define

$$\mathcal{K}_{\delta,+} \;:=\; \bigl\{x \in \mathbb{R}^I : e\cdot x > \delta|x|\bigr\}\,. \tag{4.2}$$

It follows by the proof of [14, Theorem 3.1] that a sufficient condition for BQBS stability is that (2.28) holds with $\mathcal{K}_\delta$ replaced by $\mathcal{K}_{\delta,+}$, i.e.,

$$b(x,u)\cdot\nabla\mathcal{V}_{\kappa,\beta}(x) \;\le\; c_0 + c_1\mathcal{V}_{\kappa,\beta}(x)\,\mathbb{1}_{\mathcal{K}_{\delta,+}}(x) - c_2\mathcal{V}_{\kappa,\beta}(x)\,\mathbb{1}_{\mathcal{K}^c_{\delta,+}}(x) \qquad \forall\,(x,u)\in\mathbb{R}^I\times\mathbb{U}\,. \tag{4.3}$$

As shown later in Proposition 6.2, the inequality in (4.3) is sufficient for (4.1) to hold for the n^{th} system, uniformly in $n\in\mathbb{N}$. The class of networks which satisfy (4.3), and are therefore BQBS stable, includes the following special classes:

(i) *Networks with a single dominant class*: there is only one class of jobs that can be served by more than one server pools (see Corollary 4.2 in [15]). This includes the standard "N" and "W" networks, the generalized "N" and "W" networks, and more general networks as depicted in Figure 1.

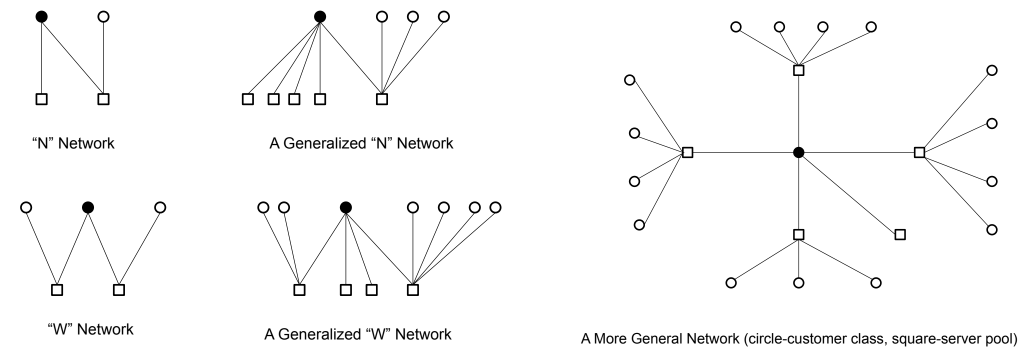

Figure 1: Examples of networks with a single dominant class

(ii) *Networks with the following parameter assumptions*:

$$\max_{i,i'\in\mathcal{I},\,j\in\mathcal{J}(i)} |\mu_{ij}-\mu_{i'j}| \;\le\; \tilde{\delta}\,\max_{i\in\mathcal{I},\,j\in\mathcal{J}}\,\{\mu_{ij}\}\,,$$

for sufficiently small $\tilde{\delta}>0$. This includes networks with pool-dependent service rates, i.e., $\mu_{ij}=\bar{\mu}_j$ for all $(i,j)\in\mathcal{E}$, as a special class. (See Corollary 4.1 in [15]).

Remark 4.1. For networks that satisfy (4.3), ergodic control problems with a running cost penalizing only the queue are well posed (e.g., we may allow $\zeta=0$ in (3.1)). This is because, in the diffusion scale, the average value of the state process is controlled by the average value of the queue, and also by the fact, as shown in Lemma 2.2, that idleness is upper bounded by some multiple of the state.

4.2. The fairness problem

In addition to ergodic control problems as in (P1)–(P2), for BQBS stable networks we can also consider constrained problems which aim at balancing idleness among the server pools, and result in a fair allocation of idle servers. Let

$$\mathcal{S}^J \;:=\; \{\theta\in(0,1)^J \,:\, e\cdot\theta = 1\}\,.$$

For the n^{th} system, we formulate this type of ergodic control problems as follows.

(P3) (*fairness*) Here we minimize the queueing cost while keeping the average idleness of the server pools balanced. Let $\theta = (\theta_1, \dots, \theta_J)^{\mathsf{T}} \in \mathcal{S}^J$ be a positive vector and let $1 \le \tilde{m} < m$. Let $\bar{J}_{\mathsf{c}} := \sum_{j \in \mathcal{J}} J_{\mathsf{c},j}$. The associated cost minimization problem becomes

$$\begin{aligned} \hat{V}^n_{\mathsf{f}}(\hat{X}^n(0)) \;&:=\; \inf_{Z^n \in \breve{\mathfrak{Z}}^n} J_{\mathsf{o}}\bigl(\hat{X}^n(0), Z^n\bigr) \\ \text{subject to} \quad J_{\mathsf{c},j}\bigl(\hat{X}^n(0), Z^n\bigr) \;&=\; \theta_j\, \bar{J}_{\mathsf{c}}\bigl(\hat{X}^n(0), Z^n\bigr)\,, \qquad j \in \mathcal{J}\,. \end{aligned}$$

For the corresponding diffusion, we have the following cost minimization problem.

(P3′) (*fairness*) The running costs r_{o}, and r_j, $j \in \mathcal{J}$, are as in (P2′). Let $\theta = (\theta_1, \dots, \theta_J)^{\mathsf{T}} \in \mathcal{S}^J$ be a positive vector, and $1 \le \tilde{m} < m$. The ergodic control problem under idleness fairness is defined as

$$\begin{aligned} \varrho^*_{\mathsf{f}}(x) \;&=\; \inf_{U \in \mathfrak{U}} J_{x,U}[r_{\mathsf{o}}] \\ \text{subject to} \quad J_{x,U}[r_j] \;&=\; \theta_j \sum_{j \in \mathcal{J}} J_{x,U}[r_j]\,, \qquad j \in \mathcal{J}\,. \end{aligned} \tag{4.4}$$

We next state an optimality result for the fairness problem (P3′). We first introduce some notation. Let

$$H_r(x,p) \;:=\; \min_{u \in \mathbb{U}} \bigl[b(x,u)\cdot p + r(x,u)\bigr]\,. \tag{4.5}$$

For $\theta = (\theta_1, \dots, \theta_J)^{\mathsf{T}} \in \mathbb{R}^J_+$ and $\lambda = (\lambda_1, \dots, \lambda_J)^{\mathsf{T}} \in \mathbb{R}^J_+$, define the running cost $h_{\theta,\lambda}$ by

$$h_{\theta,\lambda}(x,u) \;:=\; r_{\mathsf{o}}(x,u) + \sum_{j \in \mathcal{J}} \lambda_j \bigl(r_j(x,u) - \theta_j\, \bar{r}(x,u)\bigr)\,,$$

where $\bar{r} = r_1 + \cdots + r_J$. We also let

$$\mathcal{H}_{\mathsf{f}}(\theta) \;:=\; \bigl\{\pi \in \mathcal{G} : \pi(r_j) = \theta_j \pi(\bar{r})\,,\ j \in \mathcal{J}\bigr\}\,, \quad \theta \in \mathbb{R}^J_+\,.$$

The following theorem characterizes of the optimal solution of (P3′)—see Theorem 5.8 in [15] and Theorem 4.3 in [16]. The existence of solutions to the HJB equation is proved for the diffusion control problem of the "N" network, but the argument used in the proof is applicable to the general multiclass multi-pool model discussed here. The uniqueness of the solutions V_{f} follows exactly as in the proof of Theorem 3.2 in [15].

Theorem 4.1. *Suppose that the network is BQBS stable and $\gamma_i > 0$ for some $i \in \mathcal{I}$. Then the constraint in (4.4) is feasible for any positive vector $\theta = (\theta_1, \dots, \theta_J)^{\mathsf{T}} \in \mathcal{S}^J$. In addition, the following hold.*

(a) *There exists $\lambda^* \in \mathbb{R}^J_+$ such that*

$$\inf_{\pi \in \mathcal{H}(\theta)} \pi(r_{\mathsf{o}}) \;=\; \inf_{\pi \in \mathcal{G}} \pi(h_{\theta,\lambda^*}) \;=\; \varrho^*_{\mathsf{f}}\,.$$

(b) *If $\pi^* \in \mathcal{H}(\theta)$ attains the infimum of $\pi \mapsto \pi(r_{\mathsf{o}})$ in $\mathcal{H}(\theta)$, then $\pi^*(r_{\mathsf{o}}) \;=\; \pi^*(h_{\theta,\lambda^*})$, and*

$$\pi^*(h_{\theta,\lambda}) \;\le\; \pi^*(h_{\theta,\lambda^*}) \;\le\; \pi(h_{\theta,\lambda^*}) \qquad \forall\,(\pi,\lambda) \in \mathcal{G} \times \mathbb{R}^J_+\,.$$

(c) *There exists $V_{\mathsf{f}} \in \mathcal{C}^2(\mathbb{R}^I)$ satisfying*

$$\min_{u \in \mathbb{U}} \bigl[\mathcal{L}^u V_{\mathsf{f}}(x) + h_{\theta,\lambda^*}(x,u)\bigr] \;=\; \pi^*(h_{\theta,\lambda^*}) = \varrho^*_{\mathsf{f}}\,, \quad x \in \mathbb{R}^I\,.$$

(d) *A stationary Markov control $v_{\mathsf{f}} \in \mathfrak{U}_{\mathrm{SSM}}$ is optimal if and only if it satisfies*

$$H_{h_{\theta,\lambda^*}}\big(x, \nabla V_{\mathsf{f}}(x)\big) \;=\; b\big(x, v_{\mathsf{f}}(x)\big) \cdot \nabla V_{\mathsf{f}}(x) + h_{\theta,\lambda^*}\big(x, v_{\mathsf{f}}(x)\big) \quad a.e.\ in\ \mathbb{R}^I\,,$$

where $H_{h_{\theta,\lambda^*}}$ is defined in (4.5) with r replaced by h_{θ,λ^*}.

(e) *The map $\theta \mapsto \inf_{\pi \in \mathcal{H}(\theta)} \pi(r_{\mathsf{o}})$ is continuous at any feasible point $\hat{\theta}$.*

We next state the asymptotic optimality result for this class of networks.

Theorem 4.2. *For the class of networks which satisfy (4.3), and $\gamma_i > 0$ for some $i \in \mathcal{I}$, the following hold.*

(i) (*lower bound*) *there exists a positive constant C_{f} such that if a sequence $\{Z^n,\ n \in \mathbb{N}\} \subset \mathfrak{Z}$ satisfies*

$$\max_{j \in \mathcal{J}} \left| \frac{J_{\mathsf{c},j}\big(\hat{X}^n(0), Z^n\big)}{\bar{J}_{\mathsf{c}}\big(\hat{X}^n(0), Z^n\big)} - \theta_j \right| \;\le\; \epsilon \tag{4.6}$$

for some $\epsilon > 0$ and for all sufficiently large $n \in \mathbb{N}$, then

$$\liminf_{n\to\infty}\ J_{\mathsf{o}}(\hat{X}^n(0), Z^n) \;\ge\; \varrho_{\mathsf{f}}^* - C_{\mathsf{f}}\epsilon\,;$$

(ii) (*upper bound*) *for any $\epsilon > 0$, there exists a sequence $\{Z^n,\ n \in \mathbb{N}\} \subset \mathfrak{Z}$ such that (4.6) holds for all sufficiently large $n \in \mathbb{N}$, and*

$$\limsup_{n\to\infty}\ J_{\mathsf{o}}(\hat{X}^n(0), Z^n) \;\le\; \varrho_{\mathsf{f}}^* + \epsilon\,.$$

Remark 4.2. The reader will certainly notice that whereas Theorem 4.1 holds for BQBS stable networks in general, i.e., networks which satisfy (4.1), the Foster–Lyapunov condition (4.3) is assumed in Theorem 4.2 which asserts asymptotic optimality. The reason behind this, is that the corresponding BQBS stability property should hold for the diffusion-scaled processes in order to establish the lower bound, and (4.3) needs to be invoked in order to assert this property (see Proposition 6.2 in Section 6). However, (4.3) is quite natural for the models considered here.

5. A family of stabilizing policies

We introduce a class of stationary Markov scheduling policies for the general multiclass multi-pool networks which is stabilizing for the diffusion-scaled state processes in the H–W regime. Let $\mathcal{I}_\circ := \{i \in \mathcal{I} \colon \gamma_i = 0\}$. Throughout this section we fix a collection $\{N^n_{ij} \in \mathbb{N}\,,\ (i,j) \in \mathcal{E}\,, n \in \mathbb{N}\}$ which satisfies

$$\lfloor \xi^*_{ij} N^n_j \rfloor \le N^n_{ij} \le \lceil \xi^*_{ij} N^n_j \rceil\,, \qquad \text{and} \quad \sum_{i \in \mathcal{I}(j)} N^n_{ij} = N^n_j\,.$$

We also define $\bar{N}^n_i := \sum_{j \in \mathcal{I}(i)} N^n_{ij}$ for $i \in \mathcal{I}$.

Lemma 5.1. *Suppose that $\mathcal{I}_\circ \neq \mathcal{I}$. Then, given $\tilde{C}_0 > 0$, there exist a collection*

$$\{\tilde{N}^n_{ij} \in \mathbb{N}\,,\ (i,j) \in \mathcal{E}\,,\ n \in \mathbb{N}\}\,,$$

and a positive constant $\hat{C}_0$ satisfying

$$\sum_{j \in \mathcal{J}(i)} \mu^n_{ij}(\tilde{N}^n_{ij} - N^n_{ij}) \;\ge\; 2\tilde{C}_0\sqrt{n} \qquad \textit{if } i \in \mathcal{I}_\circ\,, \tag{5.1}$$

$$\big|N^n_{ij} - \tilde{N}^n_{ij}\big| \;\le\; \hat{C}_0\sqrt{n} \qquad \forall (i,j) \in \mathcal{E}\,, \tag{5.2}$$

and

$$\sum_{i\in\mathcal{I}(j)} \tilde{N}^n_{ij} = N^n_j \qquad \forall\, j\in\mathcal{J} \tag{5.3}$$

for all sufficiently large $n\in\mathbb{N}$.

Proof. Suppose, without loss of generality (WLOG), that $\mathcal{I}_\circ = \{1,\dots,I-1\}$. We claim that there exists a collection $\{\psi_{ij}\in\mathbb{R}\colon j\in\mathcal{J}(i)\,,\ i\in\mathcal{I}\}$ of real numbers and a constant $C_0>0$, satisfying $\sum_{j\in\mathcal{J}(i)}\mu_{ij}\psi_{ij} > C_0$ for all $i\in\mathcal{I}_\circ$, and

$$\sum_{i\in\mathcal{I}(j)} \psi_{ij} \;=\; 0 \qquad \forall\, j\in\mathcal{J}\,. \tag{5.4}$$

To prove the claim we use the argument of contradiction. If not, then by [38, Theorem 21.1] there exists a collection of nonnegative real numbers $\varkappa_i$, $i=1,\dots,\mathcal{I}$, such that

$$\sum_{i\in\mathcal{I}_\circ}\varkappa_i \sum_{j\in\mathcal{J}(i)} \mu_{ij}\psi_{ij} \;\le\; 0 \tag{5.5}$$

for all $\{\psi_{ij}\}$ satisfying (5.4), and $\varkappa_{\hat\imath}>0$ for some $\hat\imath\in\mathcal{I}_\circ$. Since $\mathcal{G}$ is a tree, there exists a pair of finite sequences $i_1,\dots,i_\ell$ and $j_1,\dots,j_{\ell-1}$ such that $\hat\imath = i_1$, $i_\ell = I$, and $i_k\sim j_k$, $i_{k+1}\sim j_k$ for $k=1,\dots,\ell-1$. Choosing $\psi_{i_\ell j_{\ell-1}}=-1$, $\psi_{i_{\ell-1}j_{\ell-1}}=1$, $\psi_{ij_{\ell-1}}=0$ if $i\notin\{i_{\ell-1},i_\ell\}$, and $\psi_{ij}=0$ if $j\ne j_{\ell-1}$, it follows from (5.5) that $\varkappa_{i_{\ell-1}}=0$. Thus, WLOG, we may suppose that $\varkappa_{i_2}=0$. But then replacing ψ_{i_1,j_1} with $\psi_{i_1,j_1}+C$, and ψ_{i_2,j_1} with $\psi_{i_2,j_1}-C$, the new set of numbers $\{\psi_{ij}\}$ satisfies (5.4). Therefore, by (5.5) we must have

$$\sum_{i\in\mathcal{I}_\circ}\varkappa_i \sum_{j\in\mathcal{J}(i)} \mu_{ij}\psi_{ij} + \varkappa_{i_1}\mu_{i_1j_1}C \;\le\; 0$$

for all $C\in\mathbb{R}$, which is impossible since $\varkappa_{i_1}\mu_{i_1j_1}>0$. This proves the claim.

Scaling $\{\psi_{ij}\}$ by multiplying with a constant, we may assume that

$$\sum_{j\in\mathcal{J}(i)} \mu_{ij}\psi_{ij} \;>\; 3\,\tilde{C}_0 \qquad \forall\, i\in\mathcal{I}_\circ\,. \tag{5.6}$$

For each $j\in\mathcal{J}$, if $\mathcal{I}(j)$ is a singleton, i.e., $\mathcal{I}(j)=\{i_1\}$, then we define $\tilde{N}^n_{i_1j}:=N^n_{i_1j}$. Otherwise, if $\mathcal{I}(j)=\{i_1,\dots,i_\ell\}$, then we let $\tilde{N}^n_{i_kj}:=N^n_{i_kj}+\lfloor\psi_{i_kj}\sqrt{n}\rfloor$ for $k=1,\dots,\ell-1$, and $\tilde{N}^n_{i_\ell j}:= N^n_j-\sum_{k=1}^{\ell-1}\tilde{N}^n_{i_kj}$. It is clear then that (5.1) holds for all sufficiently large n by (5.6), while (5.2) and (5.3) hold by construction. This completes the proof. □

Definition 5.1. Let $\{\tilde{N}^n_{ij}\}$ be as in Lemma 5.1, and $\tilde{N}^n_i := \sum_{j\in\mathcal{I}(i)}\tilde{N}^n_{ij}$ for $i\in\mathcal{I}$. Let $\mathcal{Z}^n$ denote the class of Markov policies z satisfying

$$z_{ij}(x) \;\le\; \tilde{N}^n_{ij} \quad \forall\, i\sim j\,, \quad \text{and} \sum_{j\in\mathcal{J}(i)} z_{ij}(x) = x_i\,, \qquad \text{if } x_i \;\le\; \tilde{N}^n_i$$

$$z_{ij}(x) \;\ge\; \tilde{N}^n_{ij} \quad \forall\, i\sim j\,, \qquad \text{if } x_i \;>\; \tilde{N}^n_i\,.$$

We refer to this class of Markov policies as *balanced saturation policies* (BSPs).

We remark that if all $\gamma_i>0$ for $i\in\mathcal{I}$, then in Definition 5.1, we may replace $\tilde{N}^n_{ij}$ and $\tilde{N}^n_i$ by N^n_{ij} and $\bar{N}^n_i$, respectively. Note that by Lemma 5.1, the quantities $\tilde{N}^n_{ij}$ and $\tilde{N}_i$ are within $\mathcal{O}(\sqrt{n})$ of the quantities N^n_{ij} and $\bar{N}_i$, which can be regarded as the 'steady-state' allocations for the n^{th} system. Thus, in the class of BSPs, if $\gamma_i > 0$ for some i, then the scheduling policy z is determined using the 'shifted' steady-state allocations $\tilde{N}^n_{ij}$ and $\tilde{N}_i$.

Note that the stabilizing policy for the 'N' network in [15] belongs to the class of BSPs. As another example, for the 'M' network, if $\gamma_i > 0$ for some $i = 1, 2$, the scheduling policy $z = z(x)$, $x \in \mathbb{Z}^2_+$, defined by

$$
\begin{aligned}
z_{11} &= x_1 \wedge N^n_1 \,, \\
z_{12} &= \begin{cases} (x_1 - N^n_1)^+ \wedge \tilde{N}^n_{12} & \text{if } x_2 \geq \tilde{N}^n_2 \\ (x_1 - N^n_1)^+ \wedge (x_2 - N^n_3)^+ & \text{otherwise}\,, \end{cases} \\
z_{22} &= \begin{cases} (x_2 - N^n_3)^+ \wedge \tilde{N}^n_{22} & \text{if } x_1 \geq \tilde{N}^n_1 \\ (x_2 - N^n_3)^+ \wedge (x_1 - N^n_1)^+ & \text{otherwise}\,, \end{cases} \\
z_{23} &= x_2 \wedge N^n_3 \,,
\end{aligned}
$$

is a BSP. If $\gamma_i > 0$ for $i = 1, 2$, then in the scheduling policy above, we can replace $\tilde{N}^n_{ij}$ and $\tilde{N}^n_i$ by N^n_{ij} and $\bar{N}^n_i$, respectively.

Using the function $\hat{x}^n$ in Definition 2.3, we can write the generator $\widehat{\mathcal{L}}^z_n$ of the diffusion-scaled state process $\hat{X}^n$ under the policy $z \in \mathcal{Z}^n$ as

$$
\widehat{\mathcal{L}}^z_n f(\hat{x}) \;=\; \mathcal{L}^z_n f\bigl(\hat{x}^n(x)\bigr)\,, \tag{5.7}
$$

where $\mathcal{L}^z_n$ is as defined in (2.7).

Recall that a $\mathbb{R}^d$-valued Markov process $\{M_t : t \geq 0\}$ is called *exponentially ergodic* if it possesses an invariant probability measure $\overline{\pi}(\mathrm{d}y)$ satisfying

$$
\lim_{t\to\infty} e^{\kappa t} \bigl\|P_t(x,\cdot) - \overline{\pi}(\cdot)\bigr\|_{\mathrm{TV}} \;=\; 0 \qquad \forall\, x \in \mathbb{R}^d\,,
$$

for some $\kappa > 0$, where $P_t(x,\cdot) := \mathbb{P}^x(M_t \in \cdot)$ denotes the transition probability of M_t, and $\|\cdot\|_{\mathrm{TV}}$ denotes the total variation norm.

Proposition 5.1. *Let $\widehat{\mathcal{L}}^z_n$ denote the generator of the diffusion-scaled state process $\hat{X}^n$ under a BSP $z \in \mathcal{Z}^n$. Let $\widetilde{\mathcal{V}}_{\epsilon,\beta}$ be as in Definition 2.6, with $\beta \in \mathbb{R}^I$ a positive vector. There exists $\epsilon > 0$, and positive constants C_0 and C_1 such that*

$$
\widehat{\mathcal{L}}^z_n\, \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \;\leq\; C_0 - C_1\, \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \qquad \forall\, \hat{x} \in \mathcal{S}^n\,, \quad \forall\, n \geq n_0\,. \tag{5.8}
$$

The process $\hat{X}^n$ is exponentially ergodic and admits a unique invariant probability measure $\hat{\pi}^n$ satisfying

$$
\lim_{t\to\infty} e^{\kappa t} \bigl\|P^n_t(x,\cdot) - \hat{\pi}^n(\cdot)\bigr\|_{\mathrm{TV}} \;=\; 0\,, \quad x \in \mathbb{R}^I\,, \tag{5.9}
$$

for any $\kappa < C_1$, where $P^n_t(x,\cdot)$ denotes the transition probability of $\hat{X}^n$.

Proof. Throughout the proof we use, without further mention, the fact that there exists a constant $\tilde{C}_0$ such that

$$
\begin{aligned}
\Bigl|\lambda^n_i - \sum_{j\in\mathcal{J}(i)} \mu^n_{ij} N^n_{ij}\Bigr| &\;\leq\; \tilde{C}_0\sqrt{n}\,, \\
\bigl|n x^*_i - \bar{N}^n_i\bigr| \;=\; \Bigl|n x^*_i - \sum_{j\in\mathcal{J}(i)} N^n_{ij}\Bigr| &\;\leq\; \tilde{C}_0\sqrt{n}
\end{aligned}
$$

for all $i \in \mathcal{I}$ and all sufficiently large $n \in \mathbb{N}$. This follows by (2.1).

Recall the collection $\{\tilde{N}^n_{ij} \in \mathbb{N}\,,\ (i,j)\in\mathcal{E}\}$ constructed in Lemma 5.1 with respect to the constant $\tilde{C}_0$ given above. Define $\breve{x}_i = \breve{x}^n_i(x) := x_i - \tilde{N}^n_i$, $\hat{\breve{x}}_i := \frac{1}{\sqrt{n}}\breve{x}^n_i$, and let

$$\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \;:=\; \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{\breve{x}})\,.$$

Using the identity

$$f(x\pm e_i) - f(x) \mp \partial_i f(x) \;=\; \int_0^1 (1-t)\,\partial_{ii} f(x\pm te_i)\,\mathrm{d}t\,, \tag{5.10}$$

we obtain

$$\Bigl|\widehat{\mathcal{V}}_{\epsilon,\beta}(x\pm e_i) - \widehat{\mathcal{V}}_{\epsilon,\beta}(x) \mp \epsilon\tfrac{\beta_i}{\sqrt{n}}\hat{\breve{x}}_i\phi_\beta(\hat{\breve{x}})\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x)\Bigr| \;\le\; \tfrac{1}{n}\epsilon^2\,\tilde{\kappa}_1\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \tag{5.11}$$

for some constant $\tilde{\kappa}_1 > 0$, and all $\epsilon\in(0,1)$, with

$$\phi_\beta(x) \;:=\; \frac{2+\|x\|^2_\beta}{\bigl(1+\|x\|^2_\beta\bigr)^{3/2}}\,.$$

Fix $n\in\mathbb{N}$. By (2.7), with

$$q_i \;=\; q_i(x_i) \;=\; x_i - \sum_{j\in\mathcal{J}(i)} z_{ij}$$

for $i\in\mathcal{I}$ (see (2.14)), and using (5.11), we obtain

$$\begin{aligned}
\mathcal{L}^z_n\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \;\le\; \epsilon\sum_{i\in\mathcal{I}}\Biggl[\lambda^n_i\Bigl(\tfrac{\beta_i}{\sqrt{n}}\hat{\breve{x}}_i\phi_\beta(\hat{\breve{x}}) + \tfrac{1}{n}\epsilon\,\tilde{\kappa}_1\Bigr) &+ \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}z_{ij}\Bigl(-\tfrac{\beta_i}{\sqrt{n}}\hat{\breve{x}}_i\phi_\beta(\hat{\breve{x}}) + \tfrac{1}{n}\epsilon\,\tilde{\kappa}_1\Bigr)\\
&+ \gamma^n_i q_i\Bigl(-\tfrac{\beta_i}{\sqrt{n}}\hat{\breve{x}}_i\phi_\beta(\hat{\breve{x}}) + \tfrac{1}{n}\epsilon\,\tilde{\kappa}_1\Bigr)\Biggr]\widehat{\mathcal{V}}_{\epsilon,\beta}(x)\\
=\; \epsilon\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x)\sum_{i\in\mathcal{I}}\Bigl(\tfrac{\beta_i}{\sqrt{n}}\,\phi_\beta(\hat{\breve{x}})\,F^{(1)}_{n,i}(x) &+ \tfrac{1}{n}\epsilon\,\tilde{\kappa}_1\,F^{(2)}_{n,i}(x)\Bigr)\,,
\end{aligned} \tag{5.12}$$

where

$$\begin{aligned}
F^{(1)}_{n,i}(x) \;&:=\; \hat{\breve{x}}_i\Biggl(\lambda^n_i - \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}z_{ij} - \gamma^n_i q_i\Biggr)\,,\\
F^{(2)}_{n,i}(x) \;&:=\; \lambda^n_i + \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}z_{ij} + \gamma^n_i q_i\,.
\end{aligned} \tag{5.13}$$

It always holds that $z_{ij}\le x_i$ and $q_i\le x_i$ for all $(i,j)\in\mathcal{E}$. By (2.1), for some constant $\tilde{\kappa}_2$ we have

$$\lambda^n_i + \sum_{j\in\mathcal{J}(i)}\bigl(\mu^n_{ij}+\gamma^n_i\bigr)\,\tilde{N}^n_i \;\le\; n\tilde{\kappa}_2 \qquad \forall\, i\in\mathcal{I}\,, \tag{5.14}$$

and all $n\in\mathbb{N}$. Thus, by (5.13) and (5.14), we obtain

$$\begin{aligned}
F^{(2)}_{n,i}(x) \;&\le\; \lambda^n_i + \sum_{j\in\mathcal{J}(i)}\mu^n_{ij}x_i + \gamma^n_i x_i\\
&=\; \lambda^n_i + \Biggl(\sum_{j\in\mathcal{J}(i)}\mu^n_{ij} + \gamma^n_i\Biggr)(\tilde{N}^n_i + \breve{x}_i)\\
&\le\; \tilde{\kappa}_2\,(n+\breve{x}_i)\,.
\end{aligned} \tag{5.15}$$

We next calculate an estimate for $F^{(1)}_{n,i}$ in (5.13). Consider any $i \in \mathcal{I}$. Define

$$\breve{z}_{ij} \;\coloneqq\; N^n_j - \sum_{i' \neq i} z_{i'j}\,.$$

We distinguish three cases.

Case A. Suppose that $x_i < \tilde{N}^n_i$. We write

$$\sum_{j \in \mathcal{J}(i)} \mu^n_{ij} z^n_{ij} \;=\; \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} \tilde{N}^n_{ij} + \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} \bigl(z^n_{ij} - \tilde{N}^n_{ij}\bigr)\,.$$

Note that $z^n_{ij} - \tilde{N}^n_{ij} \le 0$ and $\breve{x}_i \le 0$. Therefore, we have

$$-\breve{x}_i \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} \bigl(z^n_{ij} - \tilde{N}^n_{ij}\bigr) \;\le\; -\breve{x}_i \biggl(\min_{j \in \mathcal{J}(i)} \mu^n_{ij}\biggr) \bigl(x_i - \tilde{N}^n_i\bigr) \;=\; -\biggl(\min_{j \in \mathcal{J}(i)} \mu^n_{ij}\biggr) |\breve{x}_i|^2\,.$$

We also have that

$$\lambda^n_i - \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} \tilde{N}^n_{ij} \;=\; \lambda^n_i - \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} N^n_{ij} - \sum_{j \in \mathcal{J}(i)} \mu^n_{ij} (\tilde{N}^n_{ij} - N^n_{ij}) \;\le\; -\tilde{C}_0 \sqrt{n}\,. \tag{5.16}$$

Thus we obtain

$$F^{(1)}_{n,i}(x) \;\le\; -\tilde{C}_0 \sqrt{n}\, \hat{\breve{x}}_i - \sqrt{n} \biggl(\min_{j \in \mathcal{J}(i)} \mu^n_{ij}\biggr) |\hat{\breve{x}}_i|^2\,.$$

Case B. Suppose that

$$x_i \;\ge\; \tilde{N}^n_i \qquad \text{and} \qquad x_i \;\ge\; \sum_{j \in \mathcal{J}(i)} \breve{z}_{ij}\,.$$

Define

$$\breve{\zeta}_{ij} \;\coloneqq\; \breve{z}_{ij} - \tilde{N}^n_{ij}\,,$$

and note that $\breve{\zeta}_{ij} \ge 0$. Then $z_{ij} = \breve{z}_{ij} = \tilde{N}^n_{ij} + \breve{\zeta}_{ij}$.

Suppose first that $\gamma_i = 0$. Then by (5.16), we have

$$\begin{aligned}
F^{(1)}_{n,i}(x) \;&\le\; -\tilde{C}_0 \sqrt{n}\, \hat{\breve{x}}_i + \hat{\breve{x}}_i \Biggl[\lambda^n_i - \sum_{j \,\in\, \mathcal{J}(i)} \mu^n_{ij}\, \breve{z}_{ij}\Biggr] \\
&\le\; -\tilde{C}_0 \sqrt{n}\, \hat{\breve{x}}_i + \hat{\breve{x}}_i \Biggl[\lambda^n_i - \sum_{j \,\in\, \mathcal{J}(i)} \mu^n_{ij}\, (\tilde{N}^n_{ij} + \breve{\zeta}_{ij})\Biggr] \\
&\le\; -2\tilde{C}_0 \sqrt{n}\, \hat{\breve{x}}_i - \hat{\breve{x}}_i \Biggl[\sum_{j \,\in\, \mathcal{J}(i)} \mu^n_{ij}\, \breve{\zeta}_{ij}\Biggr] \\
&\le\; -2\tilde{C}_0 \sqrt{n}\, \hat{\breve{x}}_i\,.
\end{aligned}$$

Suppose now that $\gamma_i > 0$. Then by (5.16), we have

$$\begin{aligned}
F^{(1)}_{n,i}(x) \;&\le\; -\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i + \hat{\hat{x}}_i\Biggl[\lambda^n_i - \sum_{j\,\in\,\mathcal{J}(i)} \mu^n_{ij}\,\breve{z}_{ij} - \gamma^n_i\biggl(x_i - \sum_{j\,\in\,\mathcal{J}(i)} \breve{z}_{ij}\biggr)\Biggr]\\
&\le\; -\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i + \hat{\hat{x}}_i\Biggl[\lambda^n_i - \sum_{j\,\in\,\mathcal{J}(i)} \mu^n_{ij}\,(\tilde{N}^n_{ij} + \breve{\zeta}_{ij}) - \gamma^n_i\biggl(\breve{x}_i - \sum_{j\,\in\,\mathcal{J}(i)} \breve{\zeta}_{ij}\biggr)\Biggr]\\
&\le\; -2\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i - \hat{\hat{x}}_i\Biggl[\sum_{j\,\in\,\mathcal{J}(i)} \mu^n_{ij}\,\breve{\zeta}_{ij} + \gamma^n_i\biggl(\hat{\hat{x}}_i - \sum_{j\,\in\,\mathcal{J}(i)} \breve{\zeta}_{ij}\biggr)\Biggr]\\
&\le\; -2\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i - \sqrt{n}\,\Bigl(\gamma^n_i \wedge \min_{j\in\mathcal{J}(i)} \mu^n_{ij}\Bigr)\,|\hat{\hat{x}}_i|^2\,.
\end{aligned}$$

Case C. Suppose that

$$x_i \;\ge\; \tilde{N}^n_i \qquad \text{and} \qquad x_i \;<\; \sum_{j\in\mathcal{J}(i)} \breve{z}_{ij}\,.$$

Let $\hat{\jmath} \in \mathcal{J}(i)$ be arbitrary. We have

$$\begin{aligned}
F^{(1)}_{n,i}(x) \;&\le\; -\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i + \hat{\hat{x}}_i\Biggl[\lambda^n_i - \sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \mu^n_{ij}\,\breve{z}_{ij} - \mu^n_{i\hat{\jmath}}\biggl(x_i - \sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \breve{z}_{ij}\biggr)\Biggr]\\
&\le\; -\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i + \hat{\hat{x}}_i\Biggl[\lambda^n_i - \sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \mu^n_{ij}\,(\tilde{N}^n_{ij} + \breve{\zeta}_{ij}) - \mu^n_{i\hat{\jmath}}\biggl(\tilde{N}^n_{i\hat{\jmath}} + \breve{x}_i - \sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \breve{\zeta}_{ij}\biggr)\Biggr]\\
&\le\; -2\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i - \hat{\hat{x}}_i\Biggl[\sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \mu^n_{ij}\,\breve{\zeta}_{ij} + \mu^n_{i\hat{\jmath}}\biggl(\breve{x}_i - \sum_{j\,\in\,\mathcal{J}(i)\setminus\{\hat{\jmath}\}} \breve{\zeta}_{ij}\biggr)\Biggr]\\
&\le\; -2\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i - \sqrt{n}\,\Bigl(\min_{j\in\mathcal{J}(i)} \mu^n_{ij}\Bigr)\,|\hat{\hat{x}}_i|^2\,.
\end{aligned}$$

From cases A–C, we obtain

$$\begin{aligned}
F^{(1)}_{n,i}(x) \;&\le\; -2\tilde{C}_0\sqrt{n}\,\hat{\hat{x}}_i\,\mathbb{1}_{\{\breve{x}_i>0\}} - \sqrt{n}\,\biggl(\tilde{C}_0\,\hat{\hat{x}}_i + \Bigl(\min_{j\in\mathcal{J}(i)} \mu^n_{ij}\Bigr)\,|\hat{\hat{x}}_i|^2\biggr)\,\mathbb{1}_{\{\breve{x}_i\le 0\}} \qquad \text{if } \gamma_i = 0\,,\\
F^{(1)}_{n,i}(x) \;&\le\; -\sqrt{n}\,\biggl(2\tilde{C}_0\,\hat{\hat{x}}_i + \Bigl(\gamma^n_i \wedge \min_{j\in\mathcal{J}(i)} \mu^n_{ij}\Bigr)|\hat{\hat{x}}_i|^2\biggr)\,\mathbb{1}_{\{\breve{x}_i>0\}}\\
&\qquad - \sqrt{n}\,\biggl(\tilde{C}_0\,\hat{\hat{x}}_i + \Bigl(\min_{j\in\mathcal{J}(i)} \mu^n_{ij}\Bigr)\,|\hat{\hat{x}}_i|^2\biggr)\,\mathbb{1}_{\{\breve{x}_i\le 0\}} \qquad \text{if } \gamma_i > 0\,.
\end{aligned}$$

It is clear from these estimates together with (5.12) and (5.15), that, for $\varepsilon > 0$ small enough, there exist positive constants M_k, $k = 0, 1$, satisfying

$$\mathcal{L}^z_n\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \;\le\; M_0 - M_1\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \qquad \forall\, x \in \mathbb{Z}^I_+\,.$$

Define $\varsigma^n_i := \sqrt{n}x^*_i - \frac{1}{\sqrt{n}}\tilde{N}^n_i$. Then $\widehat{\mathcal{V}}_{\epsilon,\beta}(x) = \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x} + \varsigma^n)$. Thus, using the definition in (5.7) and Definition 2.3, we obtain

$$\widehat{\mathcal{L}}^z_n\,\widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x} + \varsigma^n) \;\le\; M_0 - M_1\,\widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x} + \varsigma^n)\,. \tag{5.17}$$

Since $|\varsigma_i^n| \le \hat{C}_0$ by Lemma 5.1, it is clear that (5.8) follows by (5.17).

It is standard to show that (5.8) implies (5.9). One can apply, for example, equation (3.5) in [39, Theorem 3.2], using $\Psi_1(x) = x$ and $\Psi_2(x) = 1$. This completes the proof. $\square$

The following is immediate from Proposition 5.1.

Corollary 5.1. *If $z \in \mathcal{Z}^n$ is a BSP, then for some $\varepsilon > 0$ we have*

$$\sup_{n \ge n_0} \limsup_{T \to \infty} \frac{1}{T} \, \mathbb{E}^z \left[\int_0^T \mathrm{e}^{\varepsilon \, |\hat{X}^n(s)|} \, \mathrm{d}s \right] \; < \; \infty \,, \tag{5.18}$$

and the same holds if we replace $\hat{X}^n$ with $\hat{Q}^n$ or $\hat{Y}^n$ in (5.18). In particular, the invariant probability measure of the diffusion-scaled process $\hat{X}^n(t)$ under a BSP has finite moments of any order.

Remark 5.1. A Foster–Lyapunov equation similar to (2.30) can be obtained for the diffusion-scaled state process $\hat{X}^n$ under a BSP policy. Let $\mathcal{V}_{\kappa,\beta}$ be as in Definition 2.6 and $\widehat{\mathcal{L}}_n^z$ as in Proposition 5.1. One can show, by a slight modification of the proof of Proposition 5.1, that for each $\kappa > 1$, there exist positive constants C_0 and C_1 depending only on κ and $n_0 \in \mathbb{N}$, such that for all $z \in \mathcal{Z}^n$, we have

(i) If $\gamma_i > 0$ for all $i \in \mathcal{I}$, then

$$\widehat{\mathcal{L}}_n^z \mathcal{V}_{\kappa,\beta}(\hat{x}) \;\le\; C_0 - C_1 \, \mathcal{V}_{\kappa,\beta}(\hat{x}) \qquad \forall \, \hat{x} \in \mathcal{S}^n \,, \quad \forall \, n \ge n_0 \,.$$

(ii) If $\gamma_i > 0$ for some $i \in \mathcal{I}$ (but not all i), then

$$\widehat{\mathcal{L}}_n^z \mathcal{V}_{\kappa,\beta}(\hat{x}) \;\le\; C_0 - C_1 \, \mathcal{V}_{\kappa-1,\beta}(\hat{x}) \qquad \forall \, \hat{x} \in \mathcal{S}^n \,, \quad \forall \, n \ge n_0 \,. \tag{5.19}$$

The Foster–Lyapunov property in (5.19) can be equivalently written as

$$\widehat{\mathcal{L}}_n^z \mathcal{V}_{\kappa,\beta}(\hat{x}) \;\le\; C_0 - C_1' \left(\mathcal{V}_{\kappa,\beta}(\hat{x}) \right)^{\frac{\kappa-1}{\kappa}} \qquad \forall \, \hat{x} \in \mathcal{S}^n \,, \quad \forall \, n \ge n_0 \,,$$

for some constant C_1'. Such Foster–Lyapunov properties appear in the studies on subexponential rate of convergence of Markov processes (see, e.g., [39], [40] and references therein). Thus (5.19) provides an interesting example to that rich theory. On the other hand, (5.8) with the exponential function $\widetilde{\mathcal{V}}_{\epsilon,\beta}$ is stronger, and implies exponential ergodicity of the processes $\hat{X}^n(t)$ under a BSP.

6. Ergodic Properties of the n^{th} System

6.1. Moment bounds for general multiclass multi-pool networks

Recall the moment bounds for the diffusion limit in (2.29). We prove the analogous property for the diffusion-scaled state process $\hat{X}^n$.

Proposition 6.1. *For any $\kappa \ge 1$, there exist constants $\widetilde{C}_0$ and $\widetilde{C}_1$, depending only on κ, such that*

$$\mathbb{E}^{Z^n} \left[\int_0^T |\hat{X}^n(s)|^\kappa \, \mathrm{d}s \right] \;\le\; \widetilde{C}_0 \left(T + |\hat{X}^n(0)|^\kappa \right) + \widetilde{C}_1 \, \mathbb{E}^{Z^n} \left[\int_0^T \left(1 + |\hat{Q}^n(s)| + |\hat{Y}^n(s)| \right)^\kappa \mathrm{d}s \right] \tag{6.1}$$

for all $n \in \mathbb{N}$, and for any sequence $\{Z^n \in \mathfrak{Z}^n ,\; n \in \mathbb{N}\}$.

Proof. Let $\mathcal{V}(x) := \sum_{i\in\mathcal{I}} \beta_i \mathcal{V}_i(x_i)$, $x \in \mathbb{R}^I$, where β_i, $i \in \mathcal{I}$, are positive constants to be determined later, and $\mathcal{V}_i(x_i) = |x_i|^\kappa$ when $\kappa > 1$, whereas $\mathcal{V}_i(x_i) = \frac{|x_i|^2}{\sqrt{\delta+|x_i|^2}}$ for some $\delta > 0$ when $\kappa = 1$. By applying Itô's formula on $\mathcal{V}$, we obtain from (2.12) that for $t \ge 0$,

$$\mathbb{E}\bigl[\mathcal{V}(\hat{X}^n(t))\bigr] \;=\; \mathbb{E}\bigl[\mathcal{V}(\hat{X}^n(0))\bigr] + \mathbb{E}\biggl[\int_0^t \mathcal{A}^n\mathcal{V}\bigl(\hat{X}^n(s), \hat{Z}^n(s)\bigr)\,\mathrm{d}s\biggr] + \mathbb{E}\biggl[\sum_{s\le t} \mathcal{D}\mathcal{V}(\hat{X}^n, s)\biggr], \qquad (6.2)$$

where $\mathcal{A}^n$ is given in Definition 2.4, and $\mathcal{D}\mathcal{V}(\hat{X}^n, s)$ is as in (2.20).

Let $\hat{\Theta}^n := e\cdot\hat{Q}^n \wedge e\cdot\hat{Y}^n$. Then $\hat{Q}^n = \bigl(\hat{\Theta}^n + (e\cdot\hat{X}^n)^+\bigr)\hat{u}^c$ and $\hat{Y}^n = \bigl(\hat{\Theta}^n + (e\cdot\hat{X}^n)^-\bigr)\hat{u}^s$ for some $(\hat{u}^c, \hat{u}^s) \in \mathbb{U}$ by (2.13). Also by the linearity of the map Ψ in (2.9), we obtain

$$\begin{aligned}\hat{Z}^n &= \Psi(\hat{X}^n - \hat{Q}^n, -\hat{Y}^n)\\ &= \Psi\bigl(\hat{X}^n - (e\cdot\hat{X}^n)^+\hat{u}^c, -(e\cdot\hat{X}^n)^-\hat{u}^s\bigr) - \hat{\Theta}^n\,\Psi(\hat{u}^c, \hat{u}^s)\,. \end{aligned}\qquad (6.3)$$

Define

$$\bar{\mathcal{A}}_{i,1}\bigl(x_i, \{z_{ij}\}\bigr) \;:=\; \ell_i - \sum_{i\in\mathcal{J}(i)} \mu_{ij} z_{ij} - \gamma_i\biggl(x_i - \sum_{j\in\mathcal{J}(i)} z_{ij}\biggr),$$

and

$$\bar{\mathcal{A}}\mathcal{V}(x,z) \;:=\; \sum_{i\in\mathcal{I}} \Bigl(\bar{\mathcal{A}}_{i,1}(x_i, \{z_{ij}\})\partial_i\mathcal{V}(x) + \lambda_i\,\partial_{ii}\mathcal{V}(x)\Bigr)\,. \qquad (6.4)$$

By the convergence of the parameters in (2.1), we have

$$\begin{aligned}\bigl|\bar{\mathcal{A}}_{i,1}\bigl(x_i, \{z_{ij}\}\bigr) - \mathcal{A}^n_{i,1}(x_i, \{z_{ij}\})\bigr| &\le c_1(n)\bigl(1 + \|x\|\bigr)\,,\\ \bigl|\lambda_i - \mathcal{A}^n_{i,2}(x_i, \{z_{ij}\})\bigr| &\le c_1(n)\bigl(1 + \|x\|\bigr)\,,\end{aligned}$$

for all $i \in \mathcal{I}$, for some constant $c_1(n) \searrow 0$ as $n \to \infty$. Therefore

$$\bigl|\bar{\mathcal{A}}\mathcal{V}(x,z) - \mathcal{A}^n\mathcal{V}(x,z)\bigr| \;\le\; c_1'(n)\bigl(1 + \|x\|^\kappa\bigr) \qquad (6.5)$$

for some constant $c_1'(n) \searrow 0$ as $n \to \infty$.

Recall the drift representation $b(x,u)$ in (2.26). By (6.3), we obtain that for each i,

$$\bar{\mathcal{A}}_{i,1}\bigl(\hat{X}^n_i, \{\hat{Z}^n_{ij}\}\bigr) \;=\; b_i\bigl(\hat{X}^n, (\hat{u}^c, \hat{u}^s)\bigr) + \hat{\Theta}^n \sum_{j\in\mathcal{J}(i)} \mu_{ij}\Psi_{ij}(\hat{u}^c, \hat{u}^s) - \gamma_i\,\hat{\Theta}^n\hat{u}^c_i\,. \qquad (6.6)$$

Since $-B_1$ in (2.26) is lower diagonal and Hurwitz, there exist positive constants β_i, $i \in \mathcal{I}$, such that

$$\nabla\mathcal{V}(x)\cdot B_1 x \;\ge\; c_2\mathcal{V}(x)\,,$$

for some positive constant c_2. Thus, applying Young's inequality, after some simple calculations, we obtain

$$\bar{\mathcal{A}}\mathcal{V}\bigl(\hat{X}^n, \hat{Z}^n\bigr) \;\le\; -c_3\mathcal{V}(\hat{X}^n) + c_4\bigl(1 + |e\cdot\hat{X}^n|^\kappa + (\hat{\Theta}^n)^\kappa\bigr) \qquad (6.7)$$

for some positive constants c_3 and c_4.

Concerning the last term in (6.2), we first note that by the definition of $\mathcal{V}_i$, since the jump size is of order $\frac{1}{\sqrt{n}}$, there exists a positive constant c_5 such that $\sup_{|x_i'-x_i|\le 1}|\mathcal{V}_i''(x_i')| \le c_5(1+|x_i|^{\kappa-2})$ for each $x_i \in \mathbb{R}$. Then by the Taylor remainder theorem, we obtain that for each $i \in \mathcal{I}$,

$$\Delta\mathcal{V}_i\bigl(\hat{X}^n_i(s)\bigr) - \mathcal{V}_i'\bigl(\hat{X}^n_i(s-)\bigr)\cdot\Delta\hat{X}^n_i(s) \;\le\; \frac{1}{2}\sup_{|x_i'-\hat{X}^n_i(s-)|\le 1} |\mathcal{V}_i''(x_i')|(\Delta\hat{X}^n_i(s))^2\,.$$

Thus, we have

$$\begin{aligned}
\mathbb{E}\Biggl[\sum_{s\le t}\mathcal{D}\mathcal{V}_i(\hat{X}^n,s)\Biggr] &\le \mathbb{E}\Biggl[\sum_{s\le t}c_5\Bigl(1+|\hat{X}_i^n(s-)|^{\kappa-2}\Bigr)(\Delta\hat{X}_i^n(s))^2\Biggr]\\
&= c_5\,\mathbb{E}\biggl[\int_0^t\Bigl(1+|\hat{X}_i^n(s-)|^{\kappa-2}\Bigr)\biggl(\frac{\lambda_i^n}{n}+\frac{1}{n}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n Z_{ij}^n(s)+\frac{1}{n}\gamma_i^n Q_i^n(s)\biggr)\,\mathrm{d}s\biggr]\\
&\le c_5\,\mathbb{E}\biggl[\int_0^t\Bigl(1+|\hat{X}_i^n(s-)|^{\kappa-2}\Bigr)\\
&\qquad\qquad\biggl(\frac{\lambda_i^n}{n}+\frac{1}{n}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n N_j^n+\frac{1}{n}\gamma_i^n\Bigl(\hat{\Theta}^n(s)+\bigl(e\cdot\hat{X}^n(s)\bigr)^+\Bigr)\biggr)\,\mathrm{d}s\biggr]\\
&\le \mathbb{E}\biggl[\int_0^t\biggl(\frac{c_3}{4}|\hat{X}_i^n(s)|^\kappa+c_6\Bigl(1+|e\cdot\hat{X}^n(s)|^\kappa+\bigl(\hat{\Theta}^n(s)\bigr)^\kappa\Bigr)\biggr)\,\mathrm{d}s\biggr]\,,
\end{aligned}\tag{6.8}$$

for some positive constant c_6, independent of n. In the first equality in (6.8) we use the fact that the optional martingale $[\hat{X}_i^n]$ is the sum of the squares of the jumps, and that $[\hat{X}_i^n]-\langle\hat{X}_i^n\rangle$ is a martingale, while in the last inequality we use Young's inequality.

Therefore, the assertion of the proposition follows by combining (6.2), (6.5), (6.7), and (6.8), and the inequality

$$1+|e\cdot\hat{X}^n|^\kappa+(\hat{\Theta}^n)^\kappa \;\le\; \bigl(1+|\hat{Q}^n|+|\hat{Y}^n|\bigr)^\kappa\,. \qquad \square$$

6.2. Moment bounds for BQBS stable networks

For the class of BQBS stable networks, the moment bound in (4.1) holds for the limiting diffusion X. The following proposition shows that the analogous moment bound also holds for the diffusion-scaled process $\hat{X}^n$ of this class of networks.

Proposition 6.2. *Suppose that (4.3) holds. Then Proposition 6.1 holds with (6.1) replaced by*

$$\mathbb{E}^{Z^n}\biggl[\int_0^T|\hat{X}^n(s)|^\kappa\,\mathrm{d}s\biggr] \;\le\; \widetilde{C}_0\,\bigl(T+|\hat{X}^n(0)|^\kappa\bigr)+\widetilde{C}_1\,\mathbb{E}^{Z^n}\biggl[\int_0^T\bigl(1+|\hat{Q}^n(s)|\bigr)^\kappa\,\mathrm{d}s\biggr] \tag{6.9}$$

for all $n\in\mathbb{N}$.

Proof. Recall the definition of the cone $\mathcal{K}_{\delta,+}$ in (4.2). By (4.3), (6.4), and (6.6), we obtain

$$\begin{aligned}
\bar{\mathcal{A}}\mathcal{V}\bigl(\hat{X}^n,Z^n\bigr) &\le c_0\Bigl(1+\mathcal{V}(\hat{X}^n)\mathbb{1}_{\mathcal{K}_{\delta,+}}(\hat{X}^n)+|\nabla\mathcal{V}(\hat{X}^n)|\,\hat{\Theta}^n\Bigr)-c_1\mathcal{V}(\hat{X}^n)\,\mathbb{1}_{\mathcal{K}^c_{\delta,+}}(\hat{X}^n)\\
&\le (c_0\vee c_1)\Bigl(1+\mathcal{V}(\hat{X}^n)\mathbb{1}_{\mathcal{K}_{\delta,+}}(\hat{X}^n)+|\nabla\mathcal{V}(\hat{X}^n)|\,\hat{\Theta}^n\Bigr)-c_1\mathcal{V}(\hat{X}^n)
\end{aligned}\tag{6.10}$$

for some positive constants c_0 and c_1. Since

$$\begin{aligned}
\|\hat{Q}^n\| &= \hat{\Theta}^n+(e\cdot\hat{X}^n)^+\\
&\ge \hat{\Theta}^n+\delta|\hat{X}^n|\,\mathbb{1}_{\mathcal{K}_{\delta,+}}(\hat{X}^n)\,,
\end{aligned}$$

we obtain by (6.10) that

$$\begin{aligned}
\bar{\mathcal{A}}\mathcal{V}\bigl(\hat{X}^n,Z^n\bigr) &\le (c_0\vee c_1)\Bigl(1+\|\hat{Q}^n\|+|\nabla\mathcal{V}(\hat{X}^n)|\,\|\hat{Q}^n\|\Bigr)-c_1\mathcal{V}(\hat{X}^n)\\
&\le c_0'\,\bigl(1+\mathcal{V}(\hat{Q}^n)\bigr)-c_1'\mathcal{V}(\hat{X}^n)
\end{aligned}\tag{6.11}$$

for some positive constants c_i', $i=0,1$, where the second inequality in (6.11) follows by applying Young's inequality. The rest follows as in the proof of Proposition 6.1. $\square$

As a consequence of Proposition 6.2, we also obtain the following moment bound for the idleness process.

Corollary 6.1. *If (4.3) holds, then there exist some constants $\widetilde{C}'_0 > 0$ and $\widetilde{C}'_1 > 0$ such that*

$$\mathbb{E}^{Z^n}\left[\int_0^T |\hat{Y}^n(s)|^\kappa \,\mathrm{d}s\right] \;\le\; \widetilde{C}'_0\left(T + |\hat{Y}^n(0)|^\kappa\right) + \widetilde{C}'_1\,\mathbb{E}^{Z^n}\left[\int_0^T \bigl(1 + |\hat{Q}^n(s)|\bigr)^\kappa \,\mathrm{d}s\right],$$

for all $n \in \mathbb{N}$, and for any sequence $\{Z^n \in \mathfrak{Z}^n,\ n \in \mathbb{N}\}$.

Proof. The claim follows from (6.9) and the fundamental identities $\|\hat{Q}^n\| = \hat{\Theta}^n + (e\cdot\hat{X}^n)^+$ and $\|\hat{Y}^n\| = \hat{\Theta}^n + (e\cdot\hat{X}^n)^-$. □

6.3. Convergence of mean empirical measures

For the process $\hat{X}^n$ under a scheduling policy Z^n, and with U^n as in Definition 2.5, we define the mean empirical measures

$$\Phi_T^{Z^n}(A\times B) \;:=\; \frac{1}{T}\,\mathbb{E}^{Z^n}\left[\int_0^T \mathbb{1}_{A\times B}\bigl(\hat{X}^n(t), U^n(t)\bigr)\,\mathrm{d}t\right] \tag{6.12}$$

for Borel sets $A \subset \mathbb{R}^I$ and $B \subset \mathbb{U}$. Recall Definition 2.2. The lemma which follows shows that if the long-run average first-order moment of the diffusion-scaled state process under an EJWC scheduling policy is finite, then the mean empirical measures $\Phi_T^{Z^n}$ are tight and converge to an ergodic occupation measure corresponding to some stationary stable Markov control for the limiting diffusion control problem. This property is used in the proof of the lower bounds in Theorems 3.1 and 3.2.

Lemma 6.1. *If under some sequence of scheduling policies $\{Z^n,\ n\in\mathbb{N}\} \subset \mathfrak{Z}$, we have*

$$\sup_n\ \limsup_{T\to\infty}\ \frac{1}{T}\,\mathbb{E}^{Z^n}\left[\int_0^T \bigl|\hat{X}^n(s)\bigr|\,\mathrm{d}s\right] \;<\; \infty\,, \tag{6.13}$$

then $\bigl\{\Phi_T^{Z^n} : n\in\mathbb{N},\ T>0\bigr\}$ is tight, and any limit point $\pi \in \mathcal{P}(\mathbb{R}^I\times\mathbb{U})$ of $\bigl\{\Phi_T^{Z^n}\bigr\}$ over a sequence (n_k, T_k), with $n_k\to\infty$ and $T_k\to\infty$, lies in $\mathcal{G}$.

Proof. It is clear that (6.13) implies that $\bigl\{\Phi_T^{Z^n} : n\in\mathbb{N},\ T>0\bigr\}$ is tight. For $f \in \mathcal{C}_c^\infty(\mathbb{R}^I)$, by (2.19), the definition of $\Phi_T^{Z^n}$ in (6.12), and Lemma 2.3, we obtain

$$\begin{aligned}\frac{\mathbb{E}\bigl[f(\hat{X}^n(T))\bigr] - \mathbb{E}\bigl[f(\hat{X}^n(0))\bigr]}{T} \;&=\; \int_{\mathbb{R}^I\times\mathbb{U}} \breve{\mathcal{A}}^n f(x,u)\,\Phi_T^{Z^n}(\mathrm{d}x,\mathrm{d}u)\\ &\qquad + \frac{1}{T}\,\mathbb{E}\left[\sum_{s\le T}\mathcal{D}f(\hat{X}^n, s)\right].\end{aligned} \tag{6.14}$$

Let

$$\|f\|_{\mathcal{C}^3} := \sup_{x\in\mathbb{R}^I}\Bigl(|f(x)| + \sum_{i\in\mathcal{I}}|\partial_i f(x)| + \sum_{i,i'\in\mathcal{I}}|\partial_{ii'}f(x)| + \sum_{i,i',i''\in\mathcal{I}}|\partial_{ii'i''}f(x)|\Bigr).$$

By Taylor's formula, using also the fact that the jump size is $\frac{1}{\sqrt{n}}$, we obtain

$$\bigl|\mathcal{D}f(\hat{X}^n,s)\bigr| \;\le\; \frac{\kappa\|f\|_{\mathcal{C}^3}}{\sqrt{n}}\sum_{i,i'=1}^{I}\bigl|\Delta\hat{X}_i^n(s)\Delta\hat{X}_{i'}^n(s)\bigr|\,,$$

for some constant κ that does not depend on $n \in \mathbb{N}$. On the other hand, since independent Poisson processes have no simultaneous jumps w.p.1., we have

$$\frac{1}{T}\,\mathbb{E}\biggl[\int_0^T \sum_{i,i'=1}^{I} \bigl|\Delta\hat{X}_i^n(s)\Delta\hat{X}_{i'}^n(s)\bigr|\,\mathrm{d}s\biggr] \;=\; \frac{1}{T}\,\mathbb{E}\biggl|\int_0^T \sum_{i\in\mathcal{I}}\biggl(\frac{\lambda_i^n}{n}+\frac{1}{n}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n Z_{ij}^n(s)+\frac{1}{n}\gamma_i^n Q_i^n(s)\biggr)\,\mathrm{d}s\biggr|\,,$$

and the right hand side is uniformly bounded over $n \in \mathbb{N}$ and $T > 0$ by (6.13).

Therefore, taking limits in (6.14), we obtain

$$\limsup_{n\to\infty,\,T\to\infty} \int_{\mathbb{R}^I\times\mathbb{U}} \breve{\mathcal{A}}^n f(x,u)\,\Phi_T^{Z^n}(\mathrm{d}x,\mathrm{d}u) \;=\; 0\,. \tag{6.15}$$

Let (n_k, T_k), with $n_k \to \infty$ and $T_k \to \infty$, be any sequence along which $\Phi_{T_k}^{Z^{n_k}}(\mathrm{d}x,\mathrm{d}u)$ converges to some $\pi \in \mathcal{P}(\mathbb{R}^I \times \mathbb{U})$. By (6.15), Lemma 2.5, and a standard triangle inequality, we obtain

$$\int_{\mathbb{R}^I\times\mathbb{U}} \mathcal{L}^u f(x)\,\pi(\mathrm{d}x,\mathrm{d}u) \;=\; 0\,.$$

This implies that $\pi \in \mathcal{G}$. $\square$

We introduce a canonical construction of scheduling policies which is used in the proofs of the upper bounds for asymptotic optimality. Recall Definitions 2.2 and 2.3, and $\breve{\mathfrak{X}}^n$ defined in (2.10).

Definition 6.1. Let $\varpi\colon \{x \in \mathbb{R}_+^I \colon e\cdot x \in \mathbb{Z}\} \to \mathbb{Z}_+^I$ be a measurable map defined by

$$\varpi(x) \;\coloneqq\; \Bigl(\lfloor x_1\rfloor,\dots,\lfloor x_{I-1}\rfloor,\, e\cdot x - \textstyle\sum_{i=1}^{I-1}\lfloor x_i\rfloor\Bigr)\,, \qquad x\in\mathbb{R}^I\,.$$

By abuse of notation, we denote by ϖ the similarly defined map $\varpi\colon \{x\in\mathbb{R}_+^J\colon e\cdot x\in\mathbb{Z}\}\to\mathbb{Z}_+^J$. For a precise control $v \in \mathfrak{U}_{\mathrm{SSM}}$, define the maps $q^n[v]\colon \mathbb{R}^I \to \mathbb{Z}_+^I$ and $y^n[v]\colon \mathbb{R}^I \to \mathbb{Z}_+^J$ by

$$q^n[v](\hat{x}) \;\coloneqq\; \varpi\bigl((e\cdot(\sqrt{n}\hat{x}+nx^*))^+ v^c(\hat{x})\bigr)\,, \qquad y^n[v](\hat{x}) \;\coloneqq\; \varpi\bigl((e\cdot(\sqrt{n}\hat{x}+nx^*))^- v^s(\hat{x})\bigr)\,,$$

where $\hat{x} \in \mathcal{S}^n$. Recall the definition of the linear map Ψ in (2.9). Define the Markov scheduling policy $z^n[v]$ on $\breve{\mathcal{S}}^n$ by

$$z^n[v](\hat{x}) \;\coloneqq\; \Psi\bigl(x - q^n[v](\hat{x}), N^n - y^n[v](\hat{x})\bigr)\,.$$

Corollary 6.2. *For any precise control $v \in \mathfrak{U}_{\mathrm{SSM}}$, we have*

$$e\cdot q^n[v](\hat{x}^n(x)) \wedge e\cdot y^n[v](\hat{x}^n(x)) = 0\,, \quad \textit{and} \quad z^n[v](\hat{x}^n(x)) \in \mathcal{Z}^n(x)\,,$$

for all $x \in \breve{\mathfrak{X}}^n$, i.e., the JWC condition is satisfied in $\breve{\mathfrak{X}}^n$.

Proof. This follows from Lemma 2.1 and the definition of the maps $q^n[v]$, $y^n[v]$ and $z^n[v]$. $\square$

The next lemma is used in the proof of upper bounds in Theorems 3.1 and 3.2. It shows that for any given continuous precise stationary stable Markov control, if we construct a sequence of EJWC scheduling policies as in Definition 6.1, then the corresponding mean empirical measures of the diffusion-scaled processes converge and the limit agrees with the ergodic occupation measure of the limiting diffusion corresponding to that control.

Lemma 6.2. *Let $v \in \mathfrak{U}_{\mathrm{SSM}}$ be a continuous precise control, and $\{Z^n : n \in \mathbb{N}\}$ be any sequence of admissible scheduling policies such that each Z^n agrees with the Markov scheduling policy $z^n[v]$ given in Definition 6.1 on $\sqrt{n}\breve{B}$, i.e., $Z^n(t) = z^n[v]\bigl(\hat{X}^n(t)\bigr)$ whenever $\hat{X}^n(t) \in \sqrt{n}\breve{B}$. For $\hat{x} \in \sqrt{n}\breve{B} \cap \mathcal{S}^n$, we define*

$$u^{c,n}[v](\hat{x}) \;:=\; \begin{cases} \frac{q^n[v](\hat{x})}{e\cdot q^n[v](\hat{x})} & \textit{if } e\cdot q^n[v](\hat{x}) > 0\,, \\ v^c(\hat{x}) & \textit{otherwise,} \end{cases}$$

and

$$u^{s,n}[v](\hat{x}) \;:=\; \begin{cases} \frac{y^n[v](\hat{x})}{e\cdot y^n[v](\hat{x})} & \textit{if } e\cdot y^n[v](\hat{x}) > 0\,, \\ v^s(\hat{x}) & \textit{otherwise.} \end{cases}$$

For the process X^n under the scheduling policy Z^n, define the mean empirical measures

$$\tilde{\Phi}^{Z^n}_T(A\times B) \;:=\; \frac{1}{T}\,\mathbb{E}^{Z^n}\biggl[\int_0^T \mathbb{1}_{A\times B}\bigl(\hat{X}^n(t), u^n[v]\bigl(\hat{X}^n(t)\bigr)\bigr)\,\mathrm{d}t\biggr] \tag{6.16}$$

for Borel sets $A \subset \sqrt{n}\breve{B}$ and $B \subset \mathbb{U}$. Suppose that (6.13) holds under this sequence $\{Z^n\}$. Then the ergodic occupation measure π_v of the controlled diffusion in (2.24) corresponding to v is the unique limit point in $\mathcal{P}(\mathbb{R}^d\times\mathbb{U})$ of $\tilde{\Phi}^{Z^n}_T$ as n and T tend to ∞.

Proof. The proof follows exactly as that of Lemma 7.2 in [16]. $\square$

6.4. A stability preserving property in the JWC region

If there exists a stationary Markov control under which the controlled diffusion is exponentially ergodic, then it can be shown that under the corresponding scheduling policy as constructed in Definition 6.1, the diffusion-scaled state process satisfies a Foster–Lyapunov condition of the exponential ergodicity type in the JWC region. We refer to this as the *stability preserving property in the JWC region*. This property is important to prove the upper bounds for the asymptotic optimality, and is also the reason why the spatial truncation technique works. We present this in the following Proposition.

Proposition 6.3. *Let $\widetilde{\mathcal{V}}_{\epsilon,\beta}$ be as in Definition 2.6. Suppose $v \in \mathfrak{U}_{\mathrm{SSM}}$ is such that for some positive constants c_0, c_1, and $\epsilon > 0$, and a positive vector $\beta \in \mathbb{R}^I$, it holds that*

$$\mathcal{L}^v\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \;\le\; c_0 - c_1\,\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \qquad \forall x \in \mathbb{R}^I\,. \tag{6.17}$$

Let $\hat{X}^n$ denote the diffusion-scaled state process under the scheduling policy $z^n[v]$ in Definition 6.1, and $\widehat{\mathcal{L}}_n$ denote its generator. Then, there exists $n_0 \in \mathbb{N}$ such that

$$\widehat{\mathcal{L}}_n\widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \;\le\; \hat{c}_0 - \hat{c}_1\,\widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \qquad \forall \hat{x} \in \breve{\mathcal{S}}^n \cap \sqrt{n}\breve{B}\,,$$

for some positive constants $\hat{c}_0$ and $\hat{c}_1$, and for all $n \ge n_0$.

Proof. Recall the notation $\hat{x} = \hat{x}^n(x)$ in Definition 2.3. Under the Markov scheduling policy $z^n[v]$ in Definition 6.1, for each given $x \in \mathbb{R}^I$, we define the associated diffusion-scaled quantities

$$\hat{q}^n \;=\; \hat{q}^n[v] \;:=\; \frac{q^n[v]}{\sqrt{n}}\,, \quad \hat{y}^n \;=\; \hat{y}^n[v] \;:=\; \frac{y^n[v]}{\sqrt{n}}\,, \quad \hat{z}^n \;=\; \hat{z}^n[v] \;:=\; \frac{z^n[v] - nz^*}{\sqrt{n}}\,.$$

Recall that $\widehat{\mathcal{L}}_n = \widehat{\mathcal{L}}_n^{z^n[v]}$ denotes the generator of $\hat{X}^n$ under the scheduling policy $z^n[v]$ (see (2.7) and (5.7)). Let $\widehat{\mathcal{V}}_{\epsilon,\beta}(x) := \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x})$, with $\widetilde{\mathcal{V}}_{\epsilon,\beta}$ as in Definition 2.6. Using the identity in (5.10), we obtain

$$\Bigl|\widehat{\mathcal{V}}_{\epsilon,\beta}(x \pm e_i) - \widehat{\mathcal{V}}_{\epsilon,\beta}(x) \mp \epsilon\tfrac{1}{\sqrt{n}}\,\beta_i\hat{x}_i\phi_\beta(\hat{x})\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x)\Bigr| \;\le\; \tfrac{1}{n}\epsilon^2\,\hat{\kappa}_1\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \tag{6.18}$$

for some constant $\hat{\kappa}_1 > 0$, and all $\epsilon \in (0,1)$. Thus by (2.7), (5.7), and (6.18), we obtain

$$\mathcal{L}_n^z\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x) \;\le\; \epsilon\,\widehat{\mathcal{V}}_{\epsilon,\beta}(x)\sum_{i\in\mathcal{I}}\Bigl(\beta_i\,\hat{x}_i\,\phi_\beta(\hat{x})\,G_{n,i}^{(1)}(\hat{x}) + \epsilon\,\hat{\kappa}_1\,G_{n,i}^{(2)}(\hat{x})\Bigr)\,, \tag{6.19}$$

in direct analogy to (5.12), where

$$\begin{aligned}
G_{n,i}^{(1)}(\hat{x}) &:= \ell_i^n - \sum_{j\in\mathcal{J}(i)}\mu_{ij}^n\hat{z}_{ij}^n - \gamma_i^n\hat{q}_i^n\,,\\
G_{n,i}^{(2)}(\hat{x}) &:= \frac{\lambda_i^n}{n} + \sum_{j\in\mathcal{J}(i)}\mu_{ij}^n z_{ij}^* + \frac{1}{\sqrt{n}}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n\hat{z}_{ij}^n + \frac{\gamma_i^n}{\sqrt{n}}\hat{q}_i^n\,.
\end{aligned}$$

The dependence of $G_{n,i}^{(1)}$ and $G_{n,i}^{(2)}$ on $\hat{x}$ is implicit through $z^n[v]$. By Definition 6.1, it always holds that $z_{ij}^n \le x_i$ and $q_i^n \le x_i$ for all $(i,j)\in\mathcal{E}$. Since $\frac{x_i}{\sqrt{n}} = \hat{x}_i + \sqrt{n}x_i^*$, and $\hat{z}_{ij} = \frac{x_i}{\sqrt{n}} + z_{ij}^*$, we obtain

$$\begin{aligned}
n\,G_{n,i}^{(2)}(x) &\le \lambda_i^n + \sqrt{n}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n z_{ij}^* + \sqrt{n}\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n\bigl(\hat{x}_i + \sqrt{n}(x_i^* - z_{ij}^*)\bigr) + \sqrt{n}\gamma_i^n(\hat{x}_i + \sqrt{n}x_i^*)\\
&= \lambda_i^n + n\Bigl(\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n + \gamma_i^n\Bigr)x_i^* + \sqrt{n}\Bigl(\sum_{j\in\mathcal{J}(i)}\mu_{ij}^n + \gamma_i^n\Bigr)\hat{x}_i\\
&\le \hat{\kappa}_2\bigl(n + \sqrt{n}|\hat{x}_i|\bigr)\,,
\end{aligned} \tag{6.20}$$

for some constant $\hat{\kappa}_2 > 0$, where the last inequality follows from the assumption on the parameters in (2.1).

Since the control v satisfies (6.17), we must have, for some positive constants c_0' and c_1' that

$$\sum_{i\in\mathcal{I}}\epsilon\beta_i b_i\bigl(x, v(x)\bigr)x_i\phi_\beta(x)\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \;\le\; c_0' - c_1'\widetilde{\mathcal{V}}_{\epsilon,\beta}(x) \qquad \forall x\in\mathbb{R}^I\,. \tag{6.21}$$

By Lemmas 2.3 and 2.5, we have $G_{n,i}^{(1)}(\hat{x}) \to b_i\bigl(\hat{x}, v(\hat{x})\bigr)$, uniformly over compact sets of $\mathbb{R}^I$ as $n\to\infty$. Therefore, the result follows by combining (6.19)–(6.21). □

7. Proofs of Asymptotic Optimality

We need the following lemma, which is used in the proof of the upper bound.

Lemma 7.1. *For any $\varepsilon > 0$, there exists a continuous precise control $v_\epsilon \in \mathfrak{U}_{\mathrm{SSM}}$ with the following properties:*

(a) *For some positive vector $\beta\in\mathbb{R}^I$ which does not depend on ε, and any $\kappa > 1$, we have*

$$\mathcal{L}^{v_\epsilon}\mathcal{V}_{\kappa,\beta}(x) \;\le\; c_0 - c_1\,\mathcal{V}_{\kappa,\beta}(x) \qquad \forall x\in\mathbb{R}^I \tag{7.1}$$

for some constants c_0 and c_1 depending only on κ.

(b) *With $\uppi_{v_\varepsilon}$ denoting the ergodic occupation measure corresponding to v_ε, it holds that*

$$\uppi_{v_\varepsilon}(r) \;=\; \int_{\mathbb{R}^I\times\mathbb{U}} r(x,u)\,\uppi_{v_\varepsilon}(\mathrm{d}x,\mathrm{d}u) \;<\; \varrho^* + \varepsilon\,,$$

where ϱ^* is the optimal value of problem (P1$'$).

Proof. By [15, Theorem 4.2] there exists a constant Markov control $\bar{u}$ and a positive vector $\beta\in\mathbb{R}^I$ satisfying

$$\mathcal{L}^{\bar{u}}\mathcal{V}_{\kappa,\beta}(x) \;\le\; \bar{c}_0 - \bar{c}_1\,\mathcal{V}_{\kappa,\beta}(x) \qquad \forall x\in\mathbb{R}^I \tag{7.2}$$

for all $\kappa>1$ and some constants $\bar{c}_0$ and $\bar{c}_1$. Even though not stated in that theorem, it follows from its proof that the constants $\bar{c}_0$ and $\bar{c}_1$ depend only on κ. We perturb r by adding a positive strictly convex function $f\colon\mathbb{U}\to\mathbb{R}_+$, such that the optimal value of the problem (P1$'$) with r replaced by $r+f$ is smaller than $\varrho^*+\frac{\varepsilon}{3}$. Following the proof of Theorems 4.1 and 4.2 in [14], there exists $R>0$ large enough and a stationary Markov control $\bar{v}_R$, which agrees with $\bar{u}$ on B_R^c and satisfies $\uppi_{\bar{v}_R}(r+f)<\varrho^*+\frac{2\varepsilon}{3}$. This control satisfies, for some $V_R\in\mathcal{C}^2(B_R)$,

$$\min_{u\in\mathbb{U}}\,\bigl[b(x,u)\cdot\nabla V_R + r(x,u) + f(u)\bigr] \;=\; b\bigl(x,\bar{v}_R(x)\bigr)\cdot\nabla V_R + r\bigl(x,\bar{v}_R(x)\bigr) + f\bigl(\bar{v}_R(x)\bigr) \tag{7.3}$$

for all $x\in B_R$. Since $u\mapsto\{b(x,u)\cdot p + r(x,u) + f(u)\}$ is strictly convex whenever it is not constant, it follows by (7.3) that $\bar{v}_R$ is continuous on B_R. Consider the concatenated Markov control which agrees with $\bar{v}_R$ in B_R and with $\bar{u}$ in B_R^c. As in the proof of [14, Theorem 2.2], we can employ a cut-off function to smoothen the discontinuity of this control at the concatenation boundary, and thus obtain a Markov control v_ε satisfying $\uppi_{v_\varepsilon}(r+f)<\varrho^*+\varepsilon$. Clearly then part (b) holds since f is nonnegative, while part (a) holds by (7.2) and the fact that v_ε agrees with $\bar{u}$ outside a compact set. This completes the proof. □

Concerning the constrained problem (P2$'$) and the fairness problem (P3$'$) we have the following analogous result.

Corollary 7.1. *For any $\varepsilon>0$, there exists a continuous precise control $v_\epsilon\in\mathfrak{U}_{\mathrm{SSM}}$ satisfying Lemma 7.1(a), and constants $\delta_j^\epsilon<\delta_j$, $j\in\mathcal{J}$ such that:*

(i) *In the case of problem* (P2$'$) *we have*

$$\uppi_{v_\varepsilon}(r_\circ) \;<\; \varrho_{\mathsf{c}}^* + \varepsilon\,,\quad \textit{and}\quad \uppi_{v_\varepsilon}(r_j) \;\le\; \delta_j^\epsilon\,,\quad j\in\mathcal{J}\,.$$

(ii) *In the case of problem* (P3$'$) *we have*

$$\uppi_{v_\varepsilon}(r_\circ) \;<\; \varrho_{\mathsf{f}}^* + \varepsilon\,,\quad \textit{and}\quad \uppi_{v_\varepsilon}(r_j) \;=\; \theta_j\sum_{\jmath\in\mathcal{J}}\uppi_{v_\varepsilon}(r_\jmath)\,,\qquad j\in\mathcal{J}\,.$$

Proof. Part (i) follows as in Theorem 5.7 in [15]. The proof of part (ii) is analogous. □

7.1. Proof of Theorem 3.1

Proof of the lower bound. Without loss of generality, we may suppose that for some increasing sequence $\{n_k\}\subset\mathbb{N}$, $Z^{n_k}\in\mathfrak{Z}^{n_k}$ is a collection of scheduling policies in $\mathfrak{Z}$ such that $J\bigl(\hat{X}^{n_k}(0),Z^{n_k}\bigr)$ converges to a finite value as $k\to\infty$. Denote by $\Phi_T^{n_k}$, the mean empirical measure $\Phi_T^{Z^{n_k}}$ defined in (6.12). Then by Proposition 6.1 and the definitions of the running cost $\hat{r}$ in (3.1) and $J\bigl(\hat{X}^n(0),Z^n\bigr)$ in (3.2) we obtain

$$\sup_{k\in\mathbb{N}}\;\limsup_{T\to\infty}\;\mathbb{E}^{Z^{n_k}}\biggl[\int_0^T|\hat{X}^{n_k}(s)|^m\,\mathrm{d}s\biggr] \;<\; \infty\,. \tag{7.4}$$

It is also clear by the definition of Φ^n_T and J that we can select a sequence $\{T_k\} \subset \mathbb{R}_+$, with $T_k \to \infty$, such that

$$\int_{\mathbb{R}^I \times \mathbb{U}} r(x,u)\, \Phi^{n_k}_{T_k}(\mathrm{d}x, \mathrm{d}u) \;\le\; J\bigl(\hat{X}^{n_k}(0), Z^{n_k}\bigr) + \frac{1}{k} \qquad \forall\, k \in \mathbb{N}\,. \tag{7.5}$$

By Lemma 6.1 and (7.4), $\{\Phi^{n_k}_{T_k} : k \in \mathbb{N}\}$ is tight and the limit of any converging subsequence $\{\Phi^{n'_k}_{T'_k}\}$ is in $\mathcal{G}$. Therefore it follows by (7.5) that

$$\lim_{k\to 0}\; J\bigl(\hat{X}^{n_k}(0), Z^{n_k}\bigr) \;\ge\; \inf_{\pi\in\mathcal{G}}\, \pi(r) \;=\; \varrho^*\,.$$

and this completes the proof. □

Proof of the upper bound. Recall $\mathcal{V}_{\kappa,\beta}$ in Definition 2.6. Let $\kappa = m+2$. By Lemma 7.1, there exists a continuous precise control v_ε such that the corresponding ergodic occupation measure satisfies $\pi_{v_\varepsilon}(r) \;<\; \varrho^* + \varepsilon$, and (7.1) holds.

For the n^{th} system, we construct a concatenated Markov scheduling policy $\mathring{z}^n$ as follows. Recall Definition 2.2. Inside $\breve{\mathfrak{X}}^n$, we apply the stationary policy $z^n[v_\epsilon]$ as in Definition 6.1, and outside $\breve{\mathfrak{X}}^n$, we apply some Markov scheduling policy $z \in \mathcal{Z}^n$ in Definition 5.1 that is exponentially stable. By Propositions 5.1 and 6.3 there exist positive constants $\hat{c}_0$, $\hat{c}_1$, a positive vector $\beta \in \mathbb{R}^I$, and $n_0 \in \mathbb{N}$, such that

$$\widehat{\mathcal{L}}^{\mathring{z}^n}_n\, \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \;\le\; \hat{c}_0 - \hat{c}_1\, \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x}) \qquad \forall \hat{x} \in \mathcal{S}^n\,, \quad \forall\, n \ge n_0\,. \tag{7.6}$$

This immediately implies that $\sup_{n\ge n_0} J(\hat{X}^n(0), Z^n) < \infty$. Let $\tilde{\Phi}^n_T \equiv \tilde{\Phi}^{\mathring{z}^n}_T$ be the corresponding mean empirical measures as defined in (6.16). Then the Foster–Lyapunov condition in (7.6) implies that we can choose a sequence $\{T_n\}$ such that

$$\sup_{n\ge n_0}\; \sup_{T\ge T_n}\; \int_{\mathbb{R}^I\times\mathbb{U}} \widetilde{\mathcal{V}}_{\epsilon,\beta}(\hat{x})\, \tilde{\Phi}^n_T(\mathrm{d}\hat{x}, \mathrm{d}u) \;<\; \infty\,. \tag{7.7}$$

WLOG, we assume that $T_n \to \infty$.

It is clear that $\mathring{z}^n$ can be viewed as a function of $\hat{x} \in \mathcal{S}^n$. We let $\hat{\mathring{z}}^n_{ij}(\hat{x}) \;:=\; \frac{(\mathring{z}^n_{ij}(\hat{x}) - nz^*)}{\sqrt{n}}$ as in Definition 2.3. In analogy to (2.15) and (2.16) we define

$$\begin{aligned}
\hat{\mathring{q}}^n_i(\hat{x}) \;&:=\; \hat{x}_i - \sum_{j\in\mathcal{J}(i)} \hat{\mathring{z}}^n_{ij}(\hat{x})\,, \quad \forall\, i \in \mathcal{I}\,,\\
\hat{\mathring{y}}^n_j(\hat{x}) \;&:=\; \frac{N^n_j - n\sum_{i\in\mathcal{I}(j)} z^*_{ij}}{\sqrt{n}} - \sum_{i\in\mathcal{J}(j)} \hat{\mathring{z}}^n_{ij}(\hat{x})\,, \quad \forall j \in \mathcal{J}\,.
\end{aligned}$$

The running cost $\hat{r}$ is uniformly integrable with respect to the collection $\{\tilde{\Phi}^n_T\,,\ n \in \mathbb{N}\,, T \ge 0\}$ by (7.7). Thus by Birkhoff's ergodic theorem, for any $\eta > 0$, we can choose a ball $B(\eta)$, and a sequence T_n such that

$$\left| \int_{B(\eta)\times\mathbb{U}} \hat{r}\bigl(\bigl(e\cdot\hat{\mathring{q}}^n(\hat{x})\bigr)^+ u^c,\, \bigl(e\cdot\hat{\mathring{y}}^n(\hat{x})\bigr)^+ u^s\bigr)\, \tilde{\Phi}^n_T(\mathrm{d}\hat{x}, \mathrm{d}u) - J(\hat{X}^n(0), \mathring{z}^n) \right| \;\le\; \frac{1}{n} + \eta\,, \tag{7.8}$$

for all $T \ge T_n$.

By the JWC condition on $\{\hat{x} \in \breve{\mathcal{S}}^n\}$ and Corollary 6.2, we have $\bigl(e\cdot\hat{\mathring{q}}^n(\hat{x})\bigr)^+ = (e\cdot\hat{x})^+$ and $\bigl(e\cdot\hat{\mathring{y}}^n(\hat{x})\bigr)^+ = (e\cdot\hat{x})^-$ for all $\hat{x} \in B(\eta)$, and for all large enough n. On the other hand we have

$$\sup_{(\hat{x},u)\in B(\eta)\times\mathbb{U}} \bigl|\hat{r}\bigl(\bigl(e\cdot\hat{\mathring{q}}^n(\hat{x})\bigr)^+ u^c,\, \bigl(e\cdot\hat{y}^n(\hat{\mathring{x}})\bigr)^+ u^s\bigr) - r(\hat{x}, u)\bigr| \;\xrightarrow[n\to\infty]{}\; 0\,. \tag{7.9}$$

Since v_ε is a continuous precise control then $\tilde{\Phi}^n_T$ converges to π_{v_ϵ} in $\mathcal{P}(\mathbb{R}^I \times \mathbb{U})$ as n and T tend to ∞ by Lemma 6.2. Thus, using (7.9) and a triangle inequality, we obtain

$$\int_{B(\eta)\times\mathbb{U}} \hat{r}\big((e\cdot\mathring{\hat{q}}^n(\hat{x}))^+ u^c,\, (e\cdot\mathring{\hat{y}}^n(\hat{x}))^+ u^s\big)\, \tilde{\Phi}^n_{T_n}(\mathrm{d}\hat{x},\mathrm{d}u) \xrightarrow[n\to\infty]{} \int_{B(\eta)\times\mathbb{U}} r(x,u)\,\pi_{v_\epsilon}(\mathrm{d}x,\mathrm{d}u)\,. \tag{7.10}$$

By (7.8) and (7.10) we have

$$\limsup_{n\to\infty}\; J(\hat{X}^n(0),\mathring{z}^n) \;\le\; \varrho^* + \epsilon + \eta\,.$$

Since η and ϵ are arbitrary, this completes the proof of the upper bound. □

Remark 7.1. It is clear that if the network satisfies (4.3) and $\zeta = 0$ in (3.4), then the same conclusion for the lower bound can be drawn by invoking Proposition 6.2 in the preceding proof.

7.2. Proof of Theorem 3.2

Proof of the lower bound. The proof follows by a similar argument as in the proof of the lower bound for Theorem 3.1. Let $\{Z^{n_k} \in \mathfrak{Z}^{n_k}\} \subset \mathfrak{Z}$, with $\{n_k\} \subset \mathbb{N}$ an increasing sequence, such that $J_\mathsf{o}\big(\hat{X}^{n_k}(0), Z^{n_k}\big)$ converges to a finite value. Select an increasing sequence $\{T_k\} \subset \mathbb{R}_+$ such that (7.5) holds with J replaced by $J_\circ$ and r by $r_\circ$. Following the proof of Theorem 3.1, let $\hat{\pi} \in \mathcal{P}(\mathbb{R}^I \times \mathbb{U})$ be the limit of $\Phi^{n'_k}_{T'_k}$ along some subsequence $\{n'_k, T'_k\} \subset \{n_k, T_k\}$. Recall the definition of r_j in (3.5). Since r_j is bounded below, taking limits, we obtain $\hat{\pi}(r_j) \;\le\; \delta_j$, $j \in \mathcal{J}$. Thus, by Lemmas 3.3–3.5 and Theorems 3.1–3.2 in [15], optimality implies that $\hat{\pi}(r_\mathsf{o}) \ge \varrho^*_\mathsf{c}$. As the proof of Theorem 3.1, we obtain,

$$\liminf_{k\to\infty}\; J_\mathsf{o}\big(\hat{X}^{n'_k}(0), Z^{n'_k}\big) \;\ge\; \hat{\pi}(r_\mathsf{o}) \;\ge\; \varrho^*_\mathsf{c}\,.$$

This proves the lower bound. □

Proof of the upper bound. Let $\epsilon > 0$ be given. By Corollary 7.1, there exists a continuous precise control $v_\epsilon \in \mathfrak{U}_{\mathrm{SSM}}$ and constants $\delta^\epsilon_j < \delta_j$, $j \in \mathcal{J}$, satisfying (7.1), and

$$\pi_{v_\epsilon}(r_\mathsf{o}) \le \varrho^*_\mathsf{c} + \epsilon\,, \quad \text{and} \quad \pi_{v_\epsilon}(r_j) \le \delta^\epsilon_j\,, \quad \forall j \in \mathcal{J}\,.$$

For the n^{th} system, we construct a Markov scheduling policy Z^n as in the proof of the upper bound of Theorem 3.1, by concatenating $z^n[v_\epsilon]$ and $z \in \mathfrak{Z}^n$ in Definition 5.1.

Following the proof of part (i) and choosing η small enough, i.e., $\eta < \epsilon \wedge \frac{1}{2}\min\{\delta_j - \delta^\epsilon_j,\, j \in \mathcal{J}\}$, we obtain

$$\limsup_{n\to\infty}\; J_\mathsf{o}(\hat{X}^n(0), Z^n) \;\le\; \varrho^*_\mathsf{c} + 2\epsilon\,,$$

$$\limsup_{n\to\infty}\; J_{\mathsf{c},j}\big(\hat{X}^n(0), Z^n\big) \;\le\; \frac{1}{2}(\delta_j + \delta^\epsilon_j)\,, \quad j \in \mathcal{J}\,.$$

This completes the proof of the upper bound. □

7.3. Proof of Theorem 4.2

Proof of the lower bound. The proof follows along the same lines as that of Theorem 3.2, with the only difference that we use Proposition 6.2 instead of Proposition 6.1 to assert tightness of the ergodic occupation measures. With $\hat{\uppi}$ as given in that proof, we have

$$\liminf_{k\to\infty}\; J_{\mathsf{o}}\big(\hat{X}^{n_k}(0), Z^{n_k}\big) \;\ge\; \hat{\uppi}(r_{\mathsf{o}})\,. \tag{7.11}$$

We then obtain

$$(\uptheta_j-\epsilon)\hat{\uppi}(\bar{r}) \;\le\; \hat{\uppi}(r_j) \;\le\; (\uptheta_j+\epsilon)\hat{\uppi}(\bar{r}) \qquad \forall\, j\in\mathcal{J}\,, \tag{7.12}$$

by (4.6) and the uniform integrability of

$$\frac{1}{T}\;\mathbb{E}^{Z^{n_k}}\left[\int_0^T \big(\hat{Y}^n_j(s)\big)^{\tilde{m}}\,\mathrm{d}s\right]\,, \qquad j\in\mathcal{J}\,,$$

which is asserted by Corollary 6.1. The proof is then completed using (7.11) and (7.12) and Theorem 4.1 (e). □

Proof of the upper bound. This also follows along the lines as that of Theorem 3.2. For the limiting diffusion control problem, by Theorem 5.7 and Remark 5.1 in [15], for any $\epsilon>0$, there exists a continuous precise control $v_\epsilon\in\mathfrak{U}_{\mathrm{SSM}}$ for (P3$'$) satisfying (7.1) and

$$\uppi_{v_\epsilon}(r_{\mathsf{o}}) \;\le\; \varrho_{\mathsf{f}}^* + \epsilon\,, \qquad \text{and} \quad \uppi_{v_\epsilon}(r_j) \;=\; \uptheta_j\,\uppi_{v_\epsilon}(\bar{r})\,, \quad j\in\mathcal{J}\,. \tag{7.13}$$

In addition, we have

$$\inf_{\epsilon\in(0,1)}\; \uppi_{v_\epsilon}(\bar{r}) > 0\,.$$

This follows from observing that $\{\uppi_{v_\epsilon}\,,\ \epsilon\in(0,1)\}$ is tight, and $(e\cdot x)^-$ is strictly positive on an open subset of B_1, and from applying the Harnack inequality for the density of the invariant probability measure of the limiting diffusion.

For the n^{th} system, we construct a Markov scheduling policy $\mathring{z}^n$ as in the proof of the upper bound of Theorem 3.2, and obtain

$$\begin{aligned}
\limsup_{n\to\infty}\; J_{\mathsf{o}}\big(\hat{X}^n(0), \mathring{z}^n\big) \;&\le\; \varrho_{\mathsf{f}}^* + \epsilon\,,\\
\lim_{n\to\infty}\; J_{\mathsf{c},j}\big(\hat{X}^n(0), \mathring{z}^n\big) \;&=\; \uppi_{v_\epsilon}(r_j)\,, \qquad j\in\mathcal{J}\,.
\end{aligned} \tag{7.14}$$

The result then follows by (7.13) and (7.14), thus completing the proof. □

8. Conclusion

In this work as well as in [14–16], we have studied ergodic control problems for multiclass multi-pool networks in the H–W regime under the hypothesis that at least one abandonment parameter is positive. The key technical contributions include (i) the development of a new framework of ergodic control (unconstrained and constrained) of a broad class of diffusions, (ii) the stabilization of the limiting diffusion and the diffusion-scaled state processes, and (iii) the technique to prove asymptotic optimality involving a spatial truncation and concatenation of scheduling policies that are stabilizing. The methodology and theory can be potentially used to study ergodic control of other classes of stochastic systems.

There are several open problems that remain to be solved. First, in this work, we have identified a class of BQBS stable networks as discussed in Section 4. It will be interesting to find some examples of network models in which the boundedness of the queueing process would not imply the boundedness of the state process. Second, we have studied the networks with at least one positive abandonment parameters. It remains to study the networks with no abandonment. The challenges lie in understanding the stability properties of both the limiting diffusions and the diffusion-scaled state processes. It is worth noting that the existence of a stabilizing control asserted in Theorem 2.1, which is established via the leaf elimination algorithm in [15], depends critically on the assumption that at least one abandonment parameter is positive. Although the proof of exponential ergodicity of the BSPs also relies on that assumption, this property is expected to hold with certain positive safety staffing for at least one server pool when all abandonment rates are zero.

Acknowledgements

This research was supported in part by the Army Research Office through grant W911NF-17-1-0019, and in part by the National Science Foundation through grants DMS-1715210 and DMS-1715875. In addition, the work of Ari Arapostathis was supported in part by the Office of Naval Research through grant N00014-16-1-2956.

References

[1] Z. Aksin, M. Armony, V. Mehrotra, The modern call center: a multi-disciplinary perspective on operations management research, Prod. Oper. Manag. 16 (2007) 665–688.

[2] M. Armony, S. Israelit, A. Mandelbaum, Y. N. Marmor, Y. Tseytlin, G. Yom-Tov, Patient flow in hospitals: a data-based queueing-science perspective, Stoch. Syst. 5 (1) (2015) 146–194.

[3] H. Chen, D. D. Yao, Fundamentals of queueing networks. Performance, asymptotics, and optimization, Vol. 46 of Applications of Mathematics, Springer-Verlag, New York, 2001.

[4] N. Gans, G. Koole, A. Mandelbaum, Telephone call centers: Tutorial, review and research prospects, Manuf. Serv. Oper. Manag. 5 (2003) 79–141.

[5] R. W. Hall, Patient Flow: Reducing Delay in Healthcare, Springer, New York, 2010.

[6] F. Kelly, E. Yudovina, Stochastic Networks, Cambridge University Press, 2014.

[7] A. Mandelbaum, W. A. Massey, M. I. Reiman, Strong approximations for Markovian service networks, Queueing Systems Theory Appl. 30 (1-2) (1998) 149–201.

[8] G. Pang, D. D. Yao, Heavy-traffic limits for a many-server queueing network with switchover, Adv. in Appl. Probab. 45 (3) (2013) 645–672.

[9] P. Shi, M. Chou, J. G. Dai, D. Ding, J. Sim, Models and insights for hospital inpatient operations: Time-dependent ED boarding time, Manage. Sci. 62 (1) (2016) 1–28.

[10] J. N. Tsitsikilis, K. Xu, Queueing system topologies with limited flexibility, Proceedings of ACM SIGMETRICS 41(1) (2013) 167–178.

[11] W. Whitt, Stochastic models for the design and management of customer contact centers: some research directions, Tech. rep., Columbia University (2002).

[12] W. Whitt, Stochastic-process limits. An introduction to stochastic-process limits and their application to queues, Springer Series in Operations Research, Springer-Verlag, New York, 2002.

[13] R. J. Williams, Stochastic processing networks, Annual Review of Statistics and Its Application 3 (2016) 323–345.

[14] A. Arapostathis, A. Biswas, G. Pang, Ergodic control of multi-class $M/M/N + M$ queues in the Halfin-Whitt regime, Ann. Appl. Probab. 25 (6) (2015) 3511–3570.

[15] A. Arapostathis, G. Pang, Ergodic diffusion control of multiclass multi-pool networks in the Halfin-Whitt regime, Ann. Appl. Probab. 26 (5) (2016) 3110–3153.

[16] A. Arapostathis, G. Pang, Infinite horizon average optimality of the N-network queueing model in the Halfin-Whitt regime, ArXiv e-prints 1602.03275.

[17] A. L. Stolyar, E. Yudovina, Tightness of invariant distributions of a large-scale flexible service system under a priority discipline, Stoch. Syst. 2 (2) (2012) 381–408.

[18] A. L. Stolyar, E. Yudovina, Systems with large flexible server pools: instability of “natural” load balancing, Ann. Appl. Probab. 23 (5) (2012) 2099–2183.

[19] A. L. Stolyar, Diffusion-scale tightness of invariant distributions of a large-scale flexible service system, Adv. in Appl. Probab. 47 (1) (2015) 251–269.

[20] R. Atar, A diffusion model of scheduling control in queueing systems with many servers, Ann. Appl. Probab. 15 (1B) (2005) 820–852.

[21] R. Atar, Scheduling control for queueing systems with many servers: asymptotic optimality in heavy traffic, Ann. Appl. Probab. 15 (4) (2005) 2606–2650.

[22] R. Atar, A. Mandelbaum, G. Shaikhet, Simplified control problems for multiclass many-server queueing systems, Math. Oper. Res. 34 (4) (2009) 795–812.

[23] I. Gurvich, W. Whitt, Queue-and-idleness-ratio controls in many-server service systems, Math. Oper. Res. 34 (2) (2009) 363–396.

[24] I. Gurvich, W. Whitt, Scheduling flexible servers with convex delay costs in many-server service systems, Manuf. Serv. Oper. Manag. 11 (2) (2009) 237–253.

[25] I. Gurvich, W. Whitt, Service-level differentiation in many-server service system via queue-ratio routing, Oper. Res. 58 (2) (2010) 316–328.

[26] J. G. Dai, T. Tezcan, Optimal control of parallel server systems with many servers in heavy traffic, Queueing Syst. 59 (2) (2008) 95–134.

[27] J. G. Dai, T. Tezcan, State space collapse in many-server diffusion limits of parallel server systems, Math. Oper. Res. 36 (2) (2011) 271–320.

[28] M. Armony, Dynamic routing in large-scale service systems with heterogeneous servers, Queueing Syst. 51 (3-4) (2005) 287–329.

[29] M. Armony, A. R. Ward, Fair dynamic routing in large-scale heterogeneous-server systems, Oper. Res. 58 (3) (2010) 624–637.

[30] A. R. Ward, M. Armony, Blind fair routing in large-scale service systems with heterogeneous customers and servers, Oper. Res. 61 (1) (2013) 228–243.

[31] A. Biswas, An ergodic control problem for many-sever multi-class queueing systems with help, ArXiv e-prints 1502.02779v2.

[32] D. Gamarnik, A. L. Stolyar, Multiclass multiserver queueing system in the Halfin-Whitt heavy traffic regime: asymptotics of the stationary distribution, Queueing Syst. 71 (1-2) (2012) 25–51.

[33] A. B. Dieker, X. Gao, Positive recurrence of piecewise Ornstein-Uhlenbeck processes and common quadratic Lyapunov functions, Ann. Appl. Probab. 23 (4) (2013) 1291–1317.

[34] A. Arapostathis, G. Pang, N. Sandrić, Ergodicity of Lévy-driven SDEs arising from multiclass many-server queues, ArXiv e-prints 1707.09674.

[35] A. L. Stolyar, Tightness of stationary distributions of a flexible-server system in the Halfin–Whitt asymptotic regime, Stoch. Syst. 5 (2) (2015) 239–267.

[36] R. J. Williams, On dynamic scheduling of a parallel server system with complete resource pooling, Analysis of Communication Networks: Call Centres, Traffic and Performance. Fields Inst. Commun. Amer. Math. Soc., Providence, RI. 28 (2000) 49–71.

[37] O. Kallenberg, Foundations of Modern Probability, 2nd Edition, Probability and its Applications, Springer-Verlag, New York, 2002.

[38] R. T. Rockafellar, Convex Analysis, Vol. 28 of Princeton Mathematical Series, Princeton University Press, Princeton, New Jersey, 1946.

[39] R. Douc, G. Fort, A. Guillin, Subgeometric rates of convergence of f-ergodic strong Markov processes, Stochastic Process. Appl. 119 (3) (2009) 897–923.

[40] M. Hairer, Convergence of Markov Processes, Lecture Notes, University of Warwick, available at http://www.hairer.org/notes/Convergence.pdf (2016).