

Many-Server Queues with Random Service Rates in the Halfin-Whitt Regime: A Measure-Valued Process Approach

Burak Büke¹ and Wenyi Qin¹

¹*School of Mathematics, University of Edinburgh, Edinburgh, Scotland, UK*
B.Buke@ed.ac.uk, W.Qin-3@sms.ed.ac.uk

May 10, 2019

Abstract

We consider many-server queueing systems with heterogeneous exponential servers and renewal arrivals. The service rate of each server is a random variable drawn from a given distribution. We develop a framework for analyzing the heavy traffic limit of these queues in random environment using probability measure-valued stochastic processes. We introduce the measure-valued fairness process which denotes the proportion of cumulative idleness experienced by servers whose rates fall in a Borel subset of the support of the service rates. It can be shown that these fairness processes do not converge in the usual Skorokhod- J_1 topology, hence we introduce a new notion of convergence based on shifted versions of these processes. We also introduce some useful martingales to identify limiting fairness processes under different routing policies.

1 Introduction

Many server queues have been the subject of much research due to their applicability in large scale service systems, especially in call centers. Exact analysis of many-server systems is generally intractable and one resorts to approximation methods relying on functional strong law of large numbers and functional central limit theorems. In this work, we adopt the scaling introduced in the seminal paper of Halfin and Whitt [13], where they show that it is possible to achieve high quality of service along with high utilization of resources. This is achieved by setting the number of servers to what is required to stabilize the system, generally referred as offered load plus an amount proportional to the square root of the offered load.

The conventional analysis in the literature following [13] focuses on either identical servers or servers classified in finitely many pools where servers are identical within each pool. However, in many real world applications the servers are humans who have inherently different abilities and serve with different rates which depends on the individual abilities, mood and health of the person. Hence, in the modeling process, each server requires individual attention and this generally results in either loss of Markov property and/or explosion of the state space dimensions.

A series of papers by Atar and his colleagues [2, 3, 4] tackle this individuality problem by assuming that service rates of servers in the systems are independent and identically distributed (henceforth referred as i.i.d.) random variables that are realized at the beginning and constant

through time but may not be available to the system controller. Atar [2] investigates two routing policies, namely longest-idle-server first and fastest server first, in an *ad hoc* manner. Under these two policies Atar [2] show that the many-server systems with random service rates can be approximated by one-dimensional diffusion processes with a random drift coefficient, which has the same structure for both policies, and a mean reversion coefficient which depends on the routing policy. Our setting in this paper is similar to the i.i.d. service rate setting of the work of Atar and his colleagues. The main difference is that we develop a general framework relying on probability measure-valued processes without assuming a particular routing policy and propose a generic representation for the parameters of the limiting diffusion. Technically speaking, Atar [2] adopts a Riemann-type approach in the analysis by dividing the support of the random service rates into small intervals and shows convergence as these intervals become finer. In contrast, ours is more of a Lebesgue-type analysis where we consider Borel subsets of the support of service rates and introduce a measure relying on how the total idleness is distributed among servers with different service rates. This allows us to treat the general case without assuming a specific routing policy.

Another important feature omitted in the aforementioned literature with i.i.d. service rates is the consideration of customer abandonments. Garnett et al. [9] show that abandonments plays an important role in the design of many-server systems. We believe that abandonment behavior is particularly important when service rate uncertainty is present. In the Halfin-Whitt scaling as introduced in [2], there is always a positive probability that the total service capacity is less than the offered load and hence, the queueing system is unstable. This is also reflected in the diffusion limits derived in [2] as the probability of the drift coefficient being positive, which makes a thorough steady-state analysis of the system impossible.

The key concept in our analysis is the probability measure-valued fairness process which is a left-continuous process indicating the proportion of total cumulative idleness experienced by servers belonging to a certain set. To have a well-defined fairness process, this process assumes an arbitrary probability measure as a constant value until some of the servers in the system experiences idleness and is equal to the self-normalized cumulative idleness process after the total cumulative idleness is positive. The point of singularity where the system experiences idleness for the first time raises difficulties in the analysis. This singularity point is the only possible discontinuity of the measure-valued processes and it can be shown that the fairness processes do not converge in any of the four topologies introduced by Skorokhod [25]. To overcome this issue, we define a new notion of convergence based on the shifted versions of the fairness process around this singularity point. We then relate this new notion of convergence to the convergence of Radon measure-valued cumulative idleness processes (before the normalization) in the usual Skorokhod- J_1 topology. We demonstrate the use of this result by deriving limiting fairness measures for priority-based policies such as fastest-server-first and slowest-server-first. We then show that the policy dependent parameter of the diffusion limits derived in [2] is the expected value of a random variable following the probability measure given by the fairness process.

As the next step, we focus on the problem of identifying the fairness process for specific policies. We introduce a sequence of martingales related to the total cumulative idleness experienced by the servers belonging to a given set. We show that under some mild conditions these martingales converge in probability to 0 and this provides us an alternative characterization of the weak limit of the cumulative idleness process. Then, we use this alternative characterization to show that the policy dependent parameter identified by Atar [2] for the longest-idle-server-first policy is the same for a much more general class of policies, which we call totally blind policies. We also provide

explanation for the structure of this parameter relating it to the basic properties of minimum of independent exponential random variables.

1.1 Notation

We denote the set of real numbers, set of positive real numbers and the set of positive integers as \mathbb{R}, \mathbb{R}_+ and \mathbb{N} , respectively. To simplify notation, for any $a, b \in \mathbb{R}$, we define

$$a \wedge b := \min\{a, b\} \text{ and } a \vee b := \max\{a, b\}.$$

For any separable metric space \mathbb{X} , $\mathcal{B}(\mathbb{X})$ denotes the set of all Borel subsets of \mathbb{X} . We assume that all the random elements defined in this paper lies in the generic probability space $(\Omega, \mathcal{F}, \mathbb{P})$, unless stated otherwise. We also define \mathcal{P} to denote the space of probability measures defined on $(\mathbb{R}_+, \mathcal{B}(\mathbb{R}_+))$ equipped with Prokhorov metric [6]. We denote $\mathbb{P}^n \Rightarrow \mathbb{P}$ if probability measures converge with respect to Prokhorov metric, which is equivalent to weak convergence of these probability measures. With a slight abuse of notation, we denote $X^n \Rightarrow X$, if the probability laws of the random elements X^n weakly converge to the probability law of X . We denote the σ -field generated by any random element X as $\sigma(X)$.

We denote the spaces of continuous, bounded continuous, right-continuous and left-continuous functions that map interval $[a, b] \subset \mathbb{R}$ to \mathbb{X} as $C_{\mathbb{X}}[a, b]$, $C_{\mathbb{X}}^b[a, b]$, $D_{\mathbb{X}}[a, b]$ and $G_{\mathbb{X}}[a, b]$, respectively. The space $C_{\mathbb{X}}[a, b]$ is equipped with the topology of uniform convergence, i.e., the supremum norm, and the spaces $D_{\mathbb{X}}[a, b]$ is equipped with Skorokhod- J_1 topology and d_S denote the usual Skorokhod- J_1 metric. If $b < \infty$ and $f(t) \in G_{\mathbb{X}}[a, b]$, then $g_f(t) := f(b - t) \in D_{\mathbb{X}}[a - b, 0]$ and we define the Skorokhod- J_1 metric on $G_{\mathbb{X}}[a, b]$ as $d'_S(f_1, f_2) := d_S(g_{f_1}, g_{f_2})$ for any $f_1, f_2 \in G_{\mathbb{X}}[a, b]$. Following Whitt [27], for $f_n, f \in D_{\mathbb{X}}[a, \infty)$ $f_n \rightarrow f$ in Skorokhod- J_1 topology if and only if for any $T > a$, the modifications of these functions on $D_{\mathbb{X}}[a, T]$ converge in Skorokhod- J_1 topology. The topology on $G_{\mathbb{X}}[a, \infty)$ is defined in a similar manner. For any function $f : [0, T] \rightarrow \mathbb{R}$, we also define

$$|f|_{*,T} := \sup_{0 \leq t \leq T} |f(t)|, f^+(t) := \max\{f(t), 0\} \text{ and } f^-(t) := \max\{-f(t), 0\}.$$

For functions $f : \mathbb{X} \rightarrow \mathbb{R}$ and measure ζ defined on $(\mathbb{X}, \mathcal{B}(\mathbb{X}))$, we use the inner-product notation for the integral

$$\langle\langle f, \zeta \rangle\rangle := \int_{\mathbb{X}} f(x) d\zeta(x).$$

Defining the identity function $\iota : \mathbb{R} \rightarrow \mathbb{R}$ such that $\iota(x) = x$ for all $x \in \mathbb{R}$, the expected value of a random variable following the distribution $\zeta \in \mathcal{P}$ can be denoted as $\langle\langle \iota, \zeta \rangle\rangle$.

2 Related Literature

This work is a part of the literature initiated by the seminal work of Halfin and Whitt [13], where they develop diffusion approximations by scaling arrival rates along with the number of exponential servers. This work has important practical implications as it shows both quality of service and efficiency of the system can be achieved by using the so-called square-root safety-staffing, i.e., setting the staffing level to be the number required to stabilize the incoming load plus a multiple of the square root of this number. The main line of research in this direction assumes that all the servers are identical, i.e., the service times of customers follow an exponential distribution with

a known rate independent of the chosen server. However, in most cases the servers are humans with inherently different abilities and serve customers at different rates. Gans et al. [8] provide a thorough numerical analysis of how the server heterogeneity affects the performance based on real call center data.

A common approach in modeling server heterogeneity is to assume that the servers can be grouped as “server pools” and that service rate can vary between pools, but all servers are identical within the same pool and serve with the same rate (see, e.g. [1]). We believe that there are several limitations of this approach. First, this approach does not explicitly model the inherent individual differences between servers due to human nature. Second, in general the rate of service in each pool are assumed to be known, or at the very least the pool that each server belongs to is known, which is not always possible in practice. We believe that these limitations can be remedied by modeling the service rates as random variables following the same distribution.

In the last decade, there has been some research effort to analyze queueing systems with random parameters. Following the methodology introduced by Harrison and Zeevi [14], fluid limits are used in conjunction with stochastic programming to characterize how queueing systems should be designed and controlled under parameter uncertainty. Much of this research concentrates on arrival rate uncertainty due to forecasting errors [5, 29]. Indeed, Bassamboo et al. [5] show that if the coefficient of variation of random arrival rate is greater than a certain threshold, fluid limits for the many-server systems yield more reliable approximations compared with the diffusion limits. The literature on the uncertainty related to service is relatively few. In a recent paper, Ibrahim [15] studies the problem of staffing many-server queueing systems when the actual number of servers is random using fluid limits.

Our current work is most closely related to the excellent paper by Atar [2]. In [2], Atar develops diffusion limits for many-server systems with random service rates for two routing policies, namely longest-idle-server-first and fastest-server-first in an *ad hoc* manner. To achieve this, he partitions the support of the random service rates into small intervals and approximates the system as an inverted-V system (as studied in Armony [1]) where each interval corresponds to a pool of servers. Under the stated routing policies, the inverted-V systems admit a one-dimensional diffusion limit for any finite number of pools, exhibiting a phenomenon referred as state-space collapse in the queueing literature and as the intervals become finer the inverted-V system and the system with random service rates converge together to the same limit. In a related paper, Atar and Shwartz [4] show that it is possible to efficiently operate a system with random service rates by sampling only a small portion of the service rates. Atar, Shaki and Shwartz [3] suggest a blind policy, i.e., a policy which does not require the exact knowledge of service rates, to equalize the cumulative idleness experienced by servers.

In this paper, we analyze the system with i.i.d. random service rates introduced in [2] without assuming any particular routing policy. To achieve this generality, we use stochastic processes assuming values in the space of probability measures. Measure-valued processes have been used to obtain fluid limits for single server queueing systems (e.g., [7, 10, 11]) and more recently for many-server queues by Kaspi and Ramanan [18], Kang and Ramanan [17] and Reed and Talreja [24]. In a more recent paper, Kaspi and Ramanan [19] also use measure-valued processes to come up with limits involving stochastic partial differential equations to many-server queueing systems. In this line of work, Ramanan and co-authors use the number of customers who has a certain age in the service or queue to define the measure-valued processes. Different from the aforementioned work, we use a measure-valued process to keep track of the proportion of cumulative idleness experience

by different servers in the system. This definition is motivated by the idleness-ratio used to design routing policies to control finite pool systems [12, 26].

3 Dynamics of the System Processes

In this work, we consider an infinite sequence of queueing systems, where the n th system has n servers. The arrivals at the n th system occur according to a renewal process with rate λ^n . More specifically, assuming that u_k^n s are i.i.d. random variables with mean 1 for each $k, n \in \mathbb{N}$, the number of arrivals at the n th system by time $t \geq 0$ is

$$A^n(t) = \sup \left\{ k : \sum_{i=1}^k \frac{u_i^n}{\lambda^n} \leq t \right\}.$$

Each customer can be served by any server in the system. However, the service time of each customer depends on the agent serving her/him and if the customer is served by agent k ($1 \leq k \leq n$), the service time is distributed exponentially with rate $\tilde{\mu}_k^n$. The service rates $\tilde{\mu}_k^n$ s are positive i.i.d. random variables with a general distribution $F(\mu)$ and are constant through time for each server. We assume that $\bar{\mu} = \mathbb{E}[\tilde{\mu}_k^n]$ and $\mathbb{E}[(\tilde{\mu}_k^n)^2] < \infty$. We also assume the following Halfin-Whitt type heavy traffic condition between arrival and the expected service rates:

Assumption 1. As $n \rightarrow \infty$, $n^{-1/2}(\lambda^n - n\bar{\mu}) \rightarrow \hat{\lambda}$, where $\hat{\lambda} \in \mathbb{R}$.

Each customer has an i.i.d. exponential patience time with rate γ and abandons the system if her/his service has not commenced until the patience time expires. Once the service starts, the customer stays in the system until the service is finished. The customers are served on a first-come-first-served basis and if more than one server is idle when a customer arrives, the customer is routed to a server according to a pre-specified service discipline. At this point, the only restrictions we impose on the service discipline are (i) it is non-idling, i.e., servers cannot stay idle if there is work in the system, (ii) it is non-anticipating, i.e., when the routing is done the server is chosen based on the information available so far and (iii) it is non-preemptive, i.e., once a service starts for a customer at a server, it continues at the same server until the service finishes.

For the n th system, the number of customers in the system at time t , the number of customers routed to the k th agent and the number of customers who completes service at agent k by time t are denoted $X^n(t)$, $R_k^n(t)$ and $D_k^n(t)$, respectively. The idleness process for each agent, $I_k^n(t)$, is defined to be 1 if the k th server in the n th system is idle at time t and 0 otherwise. The idleness processes play a critical role in our analysis and can be related to the routing and departure processes as

$$I_k^n(t) = I_k^n(0) - R_k^n(t) + D_k^n(t), \text{ for all } t \geq 0, 1 \leq k \leq n.$$

The routing process is a counting process and the customers are routed to server k only when it is idle, which implies

$$(1 - I_k^n(t-))dR_k^n(t) = 0, \text{ for all } t \geq 0 \text{ and } 1 \leq k \leq n.$$

The non-idling property of the policies can be expressed as

$$(X^n(t) - n)^- = \sum_{k=1}^n I_k^n(t), \forall t \geq 0, n \in \mathbb{N}. \quad (3.1)$$

Due to the exponential nature of service and patience times, the counting processes for service completions and abandonments can be modelled as time changes and/or thinings of homogeneous Poisson processes. Pang et al. [22] provide a detailed overview of pathwise construction of such processes. We take $\{S^n(t), t \geq 0\}$ to be Poisson processes with rate 1 and event epochs $\{\theta_i^n\}$. We also take a sequence of independent uniform(0,1) random variables, U_i^n , and define the sequence of random variables

$$\kappa_i^n := \min \left\{ k : \left(\sum_{j=1}^n \tilde{\mu}_j^n \right) U_i^n \leq \sum_{j=1}^k \tilde{\mu}_j^n \right\}, \text{ for all } i, n \in \mathbb{N}.$$

To construct the departure process for each server, we first consider $\{S_P^n(t) = S(\sum_{i=1}^n \tilde{\mu}_k^n t), t \geq 0\}$, a time change of the standard Poisson process, to act as the potential service completion process for the n th system, i.e., the event epochs of this process are the potential candidates for the actual service completion times. Using splitting property of Poisson processes, if $\kappa_i^n = k$, we treat the i th event of the potential service completion process S_P^n occurring at time θ_i^n as a potential service completion by server k and define $\{S_{P,k}^n(t) = \sum_{i=1}^{S_P^n(t)} \delta_k(\kappa_i^n), t \geq 0\}$ as the potential service completion process by server k in the n th system. To convert the potential times to actual service completion times, we need to apply another splitting by checking whether server k is busy right before a potential service completion time. Hence, we represent the service completion process for server k as

$$D_k^n(t) = \int_0^t (1 - I_k^n(s-)) dS_{P,k}^n(s) = \sum_{i=1}^{S_P^n(t)} (1 - I_k^n(\theta_i^n-)) \delta_k(\kappa_i^n), t \geq 0.$$

Similarly, taking $\{N^n(t), t \geq 0\}$ to be Poisson processes with rate 1, independent of other processes defined so far, we can represent the abandonment process for the n th system as

$$\left\{ N^n \left(\gamma \int_0^t (X^n(s) - n)^+ ds \right), t \geq 0 \right\}.$$

Now, we can write the dynamics of the system length process $\{X^n(t), t \geq 0\}$ as

$$X^n(t) = X^n(0) + A^n(t) - \sum_{k=1}^n D_k^n(t) - N^n \left(\gamma \int_0^t (X^n(s) - n)^+ ds \right), \forall t \geq 0, n \in \mathbb{N}.$$

We define the scaled system length process $\hat{X}^n(t) := \frac{X^n(t) - n}{\sqrt{n}}$ and after some simple algebraic manipulations, we get

$$\begin{aligned} \hat{X}^n(t) = & \hat{X}^n(0) + \frac{A^n(t) - \lambda_n}{\sqrt{n}} + \frac{\lambda_n t - n \bar{\mu} t}{\sqrt{n}} - \frac{\sum_{k=1}^n D_k^n(t) - \tilde{\mu}_k^n \int_0^t (1 - I_k^n(s)) ds}{\sqrt{n}} - \sum_{k=1}^n \frac{\tilde{\mu}_k^n \int_0^t I_k^n(s) ds}{\sqrt{n}} \\ & + \sum_{k=1}^n \frac{\tilde{\mu}_k^n t - \bar{\mu} t}{\sqrt{n}} - \frac{N^n \left(\gamma \int_0^t (X^n(s) - n)^+ ds \right) - \gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}} - \frac{\gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}}. \end{aligned} \quad (3.2)$$

We have the following assumption on the initial conditions $X^n(0)$:

Assumption 2. $\hat{X}^n(0) = \frac{X^n(0) - n}{\sqrt{n}} \Rightarrow \xi_0$.

We also define the filtrations $\mathbf{F}^n = \{\mathcal{F}_t^n : t \geq 0\}$ as

$$\mathcal{F}_t^n := \sigma \left(\{\tilde{\mu}_k^n\}_{k=1}^n, X^n(0), \{U_i^n\}_{i=1}^{S_P^n(t)}, A^n(s), S_P^n(s), \{I_k^n(s)\}_{k=1}^n, N^n \left(\gamma \int_0^s (X^n(u) - n)^+ du \right) : 0 \leq s \leq t \right)$$

for all $t \geq 0$. We are now ready to further analyze the idleness process and define the fairness process.

3.1 Analysis of the Total Idleness Process

In this section, we analyze the scaled versions of the total idleness process which we define as

$$\bar{I}^n(t) := \frac{\sum_{k=1}^n I_k^n(t)}{n}, \hat{I}_k^n(t) := \frac{I_k^n(t)}{\sqrt{n}}, \text{ for all } 1 \leq k \leq n \text{ and } \hat{I}^n(t) := \sum_{k=1}^n \hat{I}_k^n(t), t \geq 0.$$

To be able to manipulate the system equations and define the fairness process properly, we need to understand how the total idleness behaves. We prove that the total idleness experienced by time t , scales with \sqrt{n} as n increases by showing the tightness of the appropriately scaled maximum number of idle servers.

Lemma 1. *For any fixed $T > 0$ and nonidling policy, $\{|\hat{I}^n|_{*,T}\}_{n \in \mathbb{N}}$ is tight.*

Proof. We need to prove that for any $\epsilon > 0$, there exists a $K_\epsilon > 0$, such that

$$\mathbb{P}(|\hat{I}^n|_{*,T} > K_\epsilon) < \epsilon, \text{ for all } n \in \mathbb{N}.$$

Using (3.1),

$$\begin{aligned} \sum_{k=1}^n I_k^n(t) &= \left(X^n(0) - n + A^n(t) - \sum_{k=1}^n D_k^n(t) - N^n \left(\gamma \int_0^t (X(s) - n)^+ ds \right) \right)^- \\ &\leq \left(X^n(0) - n + A^n(t) - \sum_{k=1}^n S_{P,k}^n(\tilde{\mu}_k t) - N^n \left(\gamma \int_0^t (X(s) - n)^+ ds \right) \right)^-, \end{aligned}$$

which implies

$$\begin{aligned} |\hat{I}^n(t)|_{*,T} &\leq \left| \frac{X^n(0) - n}{\sqrt{n}} \right| + \left| \frac{A^n(t) - \lambda_n t}{\sqrt{n}} \right|_{*,T} + \left| \frac{(\sum_{k=1}^n S_{P,k}^n(t) - \sum_{k=1}^n \tilde{\mu}_k t)}{\sqrt{n}} \right|_{*,T} + \left| \frac{\sum_{k=1}^n \tilde{\mu}_k t - \lambda_n t}{\sqrt{n}} \right|_{*,T} \\ &\quad + \left| \frac{N^n(\gamma \int_0^t (X^n(s) - n)^+ ds) - \gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}} \right|_{*,T} + \left| \frac{\gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}} \right|_{*,T}. \end{aligned} \tag{3.3}$$

Assumptions 1 and 2, and the functional central limit theorem for renewal processes guarantees that the first four terms on the righthand side of (3.3) are tight. The fifth term corresponds to the supremum of a martingale with predictable quadratic variation

$$\left\langle \frac{N^n(\gamma \int_0^t (X^n(s) - n)^+ ds) - \gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}} \right\rangle = \frac{\gamma \int_0^t (X^n(s) - n)^+ ds}{n}.$$

Hence, using Lenglart-Rebolledo inequality (c.f. Lemma 5.7 in [22]), the tightness of the last two terms in (3.3) are spontaneously proved once we show the tightness of

$$\left| \frac{\int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}} \right|_{*,T}.$$

To do so, we define $L^n := \{0 \leq s \leq T : X^n(s) - n > 0\}$. The paths of $X^n(\cdot)$ are in $\mathbb{D}_{\mathbb{R}}[0, \infty)$ and hence, we can express L^n as the union of disjoint intervals as $L^n = \bigcup_{i=1}^{\infty} [\alpha_i^n, \beta_i^n)$.

$$\frac{\int_0^T (X^n(s) - n)^+ ds}{\sqrt{n}} = \frac{\int_{L^n} (X^n(s) - n) ds}{\sqrt{n}} = \frac{\sum_{i=1}^{\infty} \int_{\alpha_i^n}^{\beta_i^n} (X^n(s) - n) ds}{\sqrt{n}}.$$

Note that each α_i^n is the start of a busy period. Hence, if $s \in [\alpha_i^n, \beta_i^n)$

$$X^n(s) \leq X^n(\alpha_i^n) + A^n(s) - A(\alpha_i^n) - S_P^n(s) + S_P^n(\alpha_i^n).$$

Hence,

$$\begin{aligned} \frac{\int_0^T (X(s) - n)^+ ds}{\sqrt{n}} &\leq \sum_{i=1}^{\infty} \frac{X(\alpha_i^n) - n}{\sqrt{n}} + \sum_{i=1}^{\infty} \left(\frac{\int_{\alpha_i^n}^{\beta_i^n} (A(s) - A(\alpha_i^n) - \lambda^n(s - \alpha_i^n)) ds}{\sqrt{n}} \right) \\ &\quad - \sum_{i=1}^{\infty} \int_{\alpha_i^n}^{\beta_i^n} \frac{(S_P^n(s) - S_P^n(\alpha_i^n) - \sum_{k=1}^n \tilde{\mu}_k^n(s - \alpha_i^n)) + (\lambda^n - \sum_{k=1}^n \tilde{\mu}_k^n)(s - \alpha_i^n)}{\sqrt{n}} ds \\ &\leq T \frac{\max\{X^n(0) - n, 1\}}{\sqrt{n}} + 2T \sup_{0 \leq s \leq T} \left\{ \frac{|A^n(s) - \lambda^n s|}{\sqrt{n}} \right\} \\ &\quad + 2T \sup_{0 \leq s \leq T} \left\{ \frac{|S_P^n(s) - \sum_{k=1}^n \tilde{\mu}_k^n s|}{\sqrt{n}} \right\} + \frac{T|\lambda^n - \sum_{k=1}^n \tilde{\mu}_k^n|}{\sqrt{n}}. \end{aligned}$$

Using the functional central limit theorem for Poisson and renewal processes and continuous mapping theorem along with the fact that supremum is continuous in J_1 topology, the right-hand side converges weakly and hence is tight. This proves $\{|\hat{I}^n|_{*,T}\}$ is tight. \square

Corollary 1. $|\bar{I}^n|_{*,T} \xrightarrow{p} 0$ and $\left| \frac{\int_0^t (X^n(s) - n)^+ ds}{n} \right|_{*,T} \xrightarrow{p} 0$.

Proof. For any $\epsilon > 0$, $\mathbb{P}(|\bar{I}^n|_{*,T} > \epsilon) = \mathbb{P}(|\hat{I}^n|_{*,T} > \sqrt{n}\epsilon) \rightarrow 0$. The proof of the second claim is similar. \square

The proof of the following lemma is similar in nature to that of Lemma 3.1 in [2].

Lemma 2. *For any nonidling policy we have*

$$\sup_{0 \leq t \leq T} \frac{\sum_{k=1}^n \mu_k \int_0^t I_k(s) ds}{n} \xrightarrow{p} 0.$$

for any given T .

Proof. We know that $I_k(t) = \sqrt{I_k(t)}$. Using Cauchy-Schwarz inequality

$$\begin{aligned}
\sup_{0 \leq t \leq T} \frac{\sum_{k=1}^n \mu_k \int_0^t I_k(s) ds}{n} &= \frac{\sum_{k=1}^n \mu_k \int_0^T I_k(s) ds}{n} \\
&= \int_0^T \left(\sum_{k=1}^n \frac{\mu_k \sqrt{I_k(s)}}{n} \right) ds \\
&\leq \int_0^T \left(\sqrt{\sum_{k=1}^n \frac{\mu_k^2}{n}} \sqrt{\sum_{k=1}^n \frac{I_k(s)}{n}} \right) ds \\
&= \left(\sqrt{\sum_{k=1}^n \frac{\mu_k^2}{n}} \right) \int_0^T \sqrt{\frac{I(s)}{n}} ds \\
&\leq \left(\sqrt{\sum_{k=1}^n \frac{\mu_k^2}{n}} \right) T \sqrt{\frac{|I|_{*,T}}{n}}
\end{aligned}$$

Hence, our result follows from Lemma 1. \square

Our key observation is that how the total idleness is distributed among different servers plays the key role in the analysis of many server queueing systems with random service rates. In the next section, we introduce the fairness process which keeps track of the distribution of total cumulative idleness among servers.

3.2 The Fairness Process

In this section, we concentrate on the analysis of

$$\sum_{k=1}^n \frac{\tilde{\mu}_k^n \int_0^t I_k^n(s) ds}{\sqrt{n}},$$

the fifth term on the righthand side of (3.2). For any $\epsilon \geq 0$ and $n \in \mathbb{N}$ we define the random times

$$\tau_\epsilon^n := \inf \left\{ t : \int_0^t \hat{I}^n(s) ds > \epsilon \right\}.$$

Now, we are ready to define the measure-valued fairness process. Intuitively, the fairness process keeps track of how the total cumulative idleness is distributed among different servers in the system. Specifically, if the system has been idle for some time by t (i.e., $t > \tau_0^n$), the fairness process, $\eta_t^n(A)$, represents the proportion of cumulative idleness experienced by servers whose service rates, $\tilde{\mu}_k^n$ s, are in $A \in \mathcal{B}(\mathbb{R}_+)$. When the cumulative experienced idleness is 0, this proportion is not well-defined, and we choose a probability measure $\zeta \in \mathcal{P}$ as a placeholder and set $\eta_t^n(A) = \zeta(A)$ for any $t \leq \tau_0^n$ and $A \in \mathcal{B}(\mathbb{R}_+)$. We also define the ϵ -shifted version of the fairness process as the process which is equal to ζ when $t \leq \tau_\epsilon^n$ and is equal to η_t^n for $t > \tau_\epsilon^n$. Formally, we define the fairness process for the n th system and its ϵ -shifted version as

$$\eta_t^n(A) := \begin{cases} \frac{\int_0^t \sum_{k=1}^n \delta_{\mu_k}(A) I_k(s) ds}{\int_0^t \sum_{k=1}^n I_k(s) ds} & \text{if } t > \tau_0^n \\ \zeta(A) & \text{if } t \leq \tau_0^n \end{cases},$$

$$\mathcal{S}_\epsilon^n \eta_t^n(A) := \begin{cases} \eta_t^n(A) & \text{if } t > \tau_\epsilon^n \\ \zeta(A) & \text{if } t \leq \tau_\epsilon^n \end{cases} . \quad (3.4)$$

The fairness process and its shifted versions assume values in the space of probability measures and have paths in $G_{\mathcal{P}}[0, \infty)$. We are interested in the convergence of the fairness process in some sense. Unfortunately, the sequence of fairness measures is not tight in the Skorokhod- J_1 topology (or in any of the four topologies introduced in [25]). Hence, we resort to the following idea of convergence relying on the shifted processes.

Definition 1. Assume that $\tau_\epsilon^n \Rightarrow \tau_\epsilon$ for any $\epsilon \geq 0$ and define $\mathcal{S}_\epsilon \eta$ by replacing τ_ϵ^n with τ_ϵ in (3.4). We say that $\{\eta_t\}_{t \in \mathbb{R}_+}$ is the limiting fairness process if for any $\epsilon > 0$

$$\mathcal{S}_\epsilon^n \eta^n \Rightarrow \mathcal{S}_\epsilon \eta, \text{ on } G_{\mathcal{P}}[0, \infty) \text{ endowed with Skorokhod-}J_1 \text{ topology.}$$

Lemma 1 and the continuous mapping theorem guarantee that $\{\tau_\epsilon^n\}_{n \in \mathbb{N}}$ is tight. Also, if $\epsilon_1 < \epsilon_2$ and $\mathcal{S}_{\epsilon_1}^n \eta^n \Rightarrow \mathcal{S}_{\epsilon_1} \eta$, we also have $\mathcal{S}_{\epsilon_2}^n \eta^n \Rightarrow \mathcal{S}_{\epsilon_2} \eta$. Hence, rather than considering all real $\epsilon > 0$, we can concentrate on any sequence $\{\epsilon_n\}_{n \in \mathbb{N}}$ with $\epsilon_n \rightarrow 0$. For the sake of simplicity, we take $\epsilon_n = 1/n$ for the rest of this paper.

Now, our next step is to prove the tightness of the ϵ -shifted processes for any $\epsilon > 0$. Jakubowski [16] provides useful criteria to prove the tightness of measure-valued processes, by converting the problem to that of real-valued processes. This result has been used by Kaspi and Ramanan [18] in the study of measure-valued processes arising in many server queueing systems and our approach is similar. For completeness, we state Jakubowski's criteria below.

Theorem 1 (Jakubowski [16]). A sequence of stochastic processes $\{\eta^n\}_{n \in \mathbb{N}}$ taking values in $D_{\mathcal{P}}[0, T]$ is tight if and only if:

J1. (Compact Containment Condition) For each $\rho, T > 0$, there exists a compact set $\mathcal{K}_\rho \subset \mathcal{P}$ such that

$$\liminf_{n \rightarrow \infty} \mathbb{P}(\eta_t^n \in \mathcal{K}_\rho, \text{ for all } t \in [0, T]) > 1 - \rho.$$

J2. There exists a family of functions \mathbb{F} such that

- i. $H \in \mathbb{F} : \mathcal{P} \rightarrow \mathbb{R}$ and \mathbb{F} separates points in \mathcal{P} ,
- ii. \mathbb{F} is closed under addition,
- iii. For any fixed $H \in \mathbb{F}$, the sequence of functions $\{h^n(t) := H(\eta_t^n), \text{ for all } t \in \mathbb{R}\}_{n \in \mathbb{N}}$ is tight in $G_{\mathbb{R}}[0, \infty)$ endowed with usual Skorokhod- J_1 topology.

Following the example of Kaspi and Ramanan [18], we choose the family of functions

$$\mathbb{F} := \{H : \exists f \in C_{\mathbb{R}}^b[0, \infty) \text{ such that } H(\eta) = \langle f, \eta \rangle \text{ for all } \eta \in \mathcal{P}\}.$$

We impose a slightly stronger uniform integrability condition on our fairness processes.

Assumption 3. For any given $\rho, \varrho > 0$ and $T > 0$, there exists an M such that

$$\liminf_{n \rightarrow \infty} \mathbb{P} \left(\int_M^\infty \mu d\eta_t^n(\mu) < \varrho, \text{ for all } t \in [0, T] \right) > 1 - \rho.$$

This condition is slightly stronger than J1 and holds for any policy if the service rates have bounded support. Now, we are ready to prove the tightness of shifted fairness processes.

Theorem 2. *For any $\epsilon > 0$, the sequence of ϵ -shifted processes $\{\mathcal{S}_\epsilon^n \eta^n\}_{n \in \mathbb{N}}$ is tight.*

Proof. Again it is enough to show the tightness restricting our attention to all finite intervals $[0, T]$. Assumption 3 guarantees that the condition J1 holds. As η takes values in the space of probability measures, $\{\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle\}_{n \in \mathbb{N}}$ is also bounded for any $f \in C_b[0, \infty)$. To prove J2, we need to show the modulus of continuity condition holds, i.e., for any given $\rho, \varsigma > 0$, there exists a C_ϱ such that for any $\varrho < C_\varrho$

$$\liminf_{n \rightarrow \infty} \mathbb{P}(w'(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, \varrho) \geq \rho) < \varsigma,$$

where $w'(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, \varrho) := \inf w(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, (t_i, t_{i+1}])$ and

$$w(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, (t_i, t_{i+1}]) := \sup_{t_i < s, t \leq t_{i+1}} |\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle_t - \langle f, \eta^n \rangle_s|.$$

The infimum is taken over all partitions $\{t_i\}_{i \in \mathbb{N}}$ of $[0, T]$ with $|t_{i+1} - t_i| \geq \varrho$ for all i . For more details on modulus of continuity condition, we refer the reader to Chapter 3 in [6], noting that the condition we provide is modified for left-continuous functions.

For any given $t_2 > t_1 > \tau_\epsilon^n$, we have

$$\begin{aligned} & |\langle f, \mathcal{S}_\epsilon^n \eta_{t_2}^n \rangle - \langle f, \mathcal{S}_\epsilon^n \eta_{t_1}^n \rangle| \\ &= \left| \frac{\int_0^{t_2} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds}{\int_0^{t_2} \sum_{k=1}^n \hat{I}_k^n(s) ds} - \frac{\int_0^{t_1} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds}{\int_0^{t_1} \sum_{k=1}^n \hat{I}_k^n(s) ds} \right| \\ &= \left| \frac{\int_0^{t_1} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds + \int_{t_1}^{t_2} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds}{\int_0^{t_1} \sum_{k=1}^n \hat{I}_k^n(s) ds + \int_{t_1}^{t_2} \sum_{k=1}^n \hat{I}_k^n(s) ds} - \frac{\int_0^{t_1} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds}{\int_0^{t_1} \sum_{k=1}^n \hat{I}_k^n(s) ds} \right| \\ &= \left| \frac{\left(\int_{t_1}^{t_2} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds \right) \left(\int_0^{t_1} \sum_{k=1}^n \hat{I}_k^n(s) ds \right) - \left(\int_0^{t_1} \sum_{k=1}^n f(\tilde{\mu}_k^n) \hat{I}_k^n(s) ds \right) \left(\int_{t_1}^{t_2} \sum_{k=1}^n \hat{I}_k^n(s) ds \right)}{\left(\int_0^{t_2} \sum_{k=1}^n \hat{I}_k^n(s) ds \right) \left(\int_0^{t_1} \sum_{k=1}^n \hat{I}_k^n(s) ds \right)} \right| \end{aligned}$$

Hence, for any partition of $[0, T]$ such that $t_1 = 0, t_2 = \tau_\epsilon$ and $|t_{i+1} - t_i| = \varrho$ for $i \geq 2$,

$$w(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, (t_i, t_{i+1}]) \leq \frac{2\|f\|_{*,\infty} \|I\|_{*,\infty} \varrho}{\epsilon}$$

By definition $\tau_\epsilon > \epsilon/\|I\|_{*,\infty}$ and

$$\mathbb{P}(w'(\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle, \varrho) \geq \rho) \leq \mathbb{P}\left(\|I\|_{*,\infty} > \max\left\{\frac{\rho\epsilon}{2\|f\|_{*,\infty}\varrho}, \frac{\epsilon}{\varrho}\right\}\right).$$

The tightness of $|I|_{*,\infty}$ guarantees that this probability approaches 0 as $\varrho \rightarrow 0$, which proves our theorem. \square

In our analysis of the system length process, we are mainly concerned with the process $\langle\langle \iota, \eta^n \rangle\rangle$ which is equal to the expected value of a random variable following the probability distribution η_t^n at any time t and our next result shows that tightness of the shifted versions of these processes is implied by Theorem 2.

Corollary 2. *For any $\epsilon > 0$, the set of real-valued processes $\{\langle\langle \iota, \mathcal{S}_\epsilon^n \eta^n \rangle\rangle\}_{n \in \mathbb{N}}$ is tight.*

Proof. The compact containment condition is implied by Assumption 3 as for any given $\rho, \varrho > 0$ and $T > 0$

$$\liminf_{n \rightarrow \infty} \mathbb{P}(\langle\langle \iota, \mathcal{S}_\epsilon^n \eta_t^n \rangle\rangle < M + \varrho, \text{ for all } t \in [0, T]) > 1 - \rho.$$

We now prove the modulus of continuity condition, we first define the truncated identity function as $\iota_M(\mu) = \mu \delta_{\{\mu \leq M\}}(\mu) + M \delta_{\{\mu > M\}}(\mu)$. Using that for any $t_1, t_2 \geq 0$,

$$\langle\langle \iota, \mathcal{S}_\epsilon^n \eta_{t_2}^n \rangle\rangle - \langle\langle \iota, \mathcal{S}_\epsilon^n \eta_{t_1}^n \rangle\rangle \leq \langle\langle \iota_M, \mathcal{S}_\epsilon^n \eta_{t_2}^n \rangle\rangle - \langle\langle \iota_M, \mathcal{S}_\epsilon^n \eta_{t_1}^n \rangle\rangle + \int_M^\infty \mu d\mathcal{S}_\epsilon^n \eta_{t_2}^n + \int_M^\infty \mu d\mathcal{S}_\epsilon^n \eta_{t_1}^n.$$

we write

$$\mathbb{P}(w'(\langle\langle \iota, \mathcal{S}_\epsilon^n \eta^n \rangle\rangle, \varrho) \geq \rho) \leq \mathbb{P}(w'(\langle\langle \iota_M, \mathcal{S}_\epsilon^n \eta^n \rangle\rangle, \varrho) \geq \rho - 2\varsigma) + \mathbb{P}\left(\int_M^\infty \mu d\eta_t^n(\mu) > \varsigma, \text{ for some } t \in [0, T]\right).$$

For any $\rho > 0$ we can find Δ, ϱ and M such that the righthand side is less than ς for all n , and our result is proved. \square

Choosing $\epsilon_n = 1/n$ and using a diagonal argument, Theorem 2 implies that any sequence of $\{\eta^n\}_{n \in \mathbb{N}}$ has a subsequence where shifted fairness processes weakly converges for all $\epsilon > 0$. However, doing the analysis for each $\epsilon > 0$ can be tedious for specific applications. We define the measure-valued scaled total cumulative idleness process as

$$\hat{C}^n(t, A) := \sum_{k=1}^n \delta_{\bar{\mu}_k^n}(A) \int_0^t \hat{I}_k^n(s) ds, \text{ for all } t \geq 0, A \in \mathcal{B}(\mathbb{R}_+), n \in \mathbb{N}.$$

We now show that it is possible to identify the limiting fairness process, by examining $\hat{C}^n(t, A)$ rather than the ϵ -shifted versions of the fairness process.

Lemma 3. *Suppose for all $A \in \mathcal{B}(\mathbb{R}_+)$ and $T > 0$,*

$$\lim_{n \rightarrow \infty} \sup_{0 \leq t \leq T} \left| \hat{C}^n(t, A) - \hat{C}(t, A) \right| \xrightarrow{p} 0. \quad (3.5)$$

Then

$$\eta_t(A) = \begin{cases} \frac{\hat{C}(t, A)}{\hat{C}(t, \mathbb{R}_+)} & \text{if } \hat{C}(t, \mathbb{R}_+) > 0 \\ \zeta(A) & \text{otherwise} \end{cases}$$

is the limiting fairness process.

Proof. It is clear that $\{\eta(t, \cdot)\}$ is a probability measure valued process. For any Borel set $A \in \mathcal{B}(\mathbb{R}_+)$ processes $\{\hat{C}^n(t, A), t \geq 0\}$ are continuous processes and converge in probability to $\{\hat{C}(t, A), t \geq 0\}$ in $\mathbb{C}_{\mathbb{R}_+}[0, \infty)$. The stopping times $\tau_\epsilon^n = \inf\{t > 0 : \hat{C}^n(t, \mathbb{R}_+) > \epsilon\}$ and can be thought of as a continuous function of $\hat{C}^n(\cdot, \mathbb{R}_+)$. Using the continuous mapping theorem,

$$\tau_\epsilon^n \xrightarrow{P} \tau_\epsilon = \inf\{t > 0 : \hat{C}(t, \mathbb{R}_+) > 0\}.$$

For any fixed $\epsilon > 0$, define

$$\Lambda^n(t) = \begin{cases} \frac{\tau_\epsilon}{\tau_\epsilon^n} & \text{if } 0 \leq t \leq \tau_\epsilon^n \\ \frac{\tau_\epsilon}{T - \tau_\epsilon^n} t + \frac{T(\tau_\epsilon - \tau_\epsilon^n)}{T - \tau_\epsilon^n} & \text{if } \tau_\epsilon^n < t \leq T. \end{cases}$$

Clearly, $|\Lambda^n(t) - t| \xrightarrow{P} 0$. We now prove for any $A \in \mathcal{B}(\mathbb{R}_+)$

$$\lim_{n \rightarrow \infty} |\eta^n(t, A) - \eta(\Lambda^n(t), A)| \xrightarrow{P} 0, \text{ in } G_{\mathbb{R}}[0, T]. \quad (3.6)$$

For $t \leq \tau_\epsilon^n$, we have $|\eta^n(t, A) - \eta(\Lambda^n(t), A)| = 0$ for all A . $\Lambda^n(t)$ is defined so that $\tau_\epsilon > \tau_\epsilon^n$ if and only if $t < \Lambda^n(t)$ for all t . For $t \geq \tau_\epsilon^n$

$$\begin{aligned} |\eta^n(t, A) - \eta(\Lambda^n(t), A)| &= \delta(\tau_\epsilon > \tau_\epsilon^n) |\eta^n(t, A) - \eta(\Lambda^n(t), A)| + \delta(\tau_\epsilon \leq \tau_\epsilon^n) |\eta^n(t, A) - \eta(\Lambda^n(t), A)| \\ &\leq \delta(\tau_\epsilon \leq \tau_\epsilon^n) (|\eta^n(t, A) - \eta(t, A)| + |\eta(\Lambda^n(t), A) - \eta(t, A)|) \\ &\quad + \delta(\tau_\epsilon > \tau_\epsilon^n) (|\eta^n(\Lambda^n(t), A) - \eta(\Lambda^n(t), A)| + |\eta^n(\Lambda^n(t), A) - \eta^n(t, A)|) \\ &= \delta(\tau_\epsilon \leq \tau_\epsilon^n) \left(\left| \frac{\hat{C}^n(t, A)}{\hat{C}^n(t, \mathbb{R}_+)} - \frac{\hat{C}(t, A)}{\hat{C}(t, \mathbb{R}_+)} \right| + \left| \frac{\hat{C}^n(\Lambda^n(t), A)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)} - \frac{\hat{C}^n(t, A)}{\hat{C}^n(t, \mathbb{R}_+)} \right| \right) \\ &\quad + \delta(\tau_\epsilon > \tau_\epsilon^n) \left(\left| \frac{\hat{C}^n(\Lambda^n(t), A)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)} - \frac{\hat{C}(\Lambda^n(t), A)}{\hat{C}(\Lambda^n(t), \mathbb{R}_+)} \right| + \left| \frac{\hat{C}^n(\Lambda^n(t), A)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)} - \frac{\hat{C}^n(t, A)}{\hat{C}^n(t, \mathbb{R}_+)} \right| \right). \end{aligned}$$

We think of the righthand side of the last equality as four summands and if we can show converge of each summand to 0 in probability, then we can conclude that for all A , $\eta^n(\cdot, A) \xrightarrow{P} \eta(\cdot, A)$ in $G_{\mathbb{R}_+}[0, T]$. For the first summand, we have

$$\begin{aligned} \left| \frac{\hat{C}^n(t, A)}{\hat{C}^n(t, \mathbb{R}_+)} - \frac{\hat{C}(t, A)}{\hat{C}(t, \mathbb{R}_+)} \right| &= \left| \frac{\hat{C}(t, \mathbb{R}_+)(\hat{C}^n(t, A) - \hat{C}(t, A)) + \hat{C}(t, A)(\hat{C}^n(t, \mathbb{R}_+) - \hat{C}(t, \mathbb{R}_+))}{\hat{C}(t, \mathbb{R}_+)\hat{C}^n(t, \mathbb{R}_+)} \right| \\ &\leq \frac{\hat{I}^n(t)T \left(\left| \hat{C}^n(t, A) - \hat{C}(t, A) \right| + \left| \hat{C}^n(t, \mathbb{R}_+) - \hat{C}(t, \mathbb{R}_+) \right| \right)}{\epsilon^2}, \end{aligned}$$

and for the second summand,

$$\begin{aligned} \left| \frac{\hat{C}^n(\Lambda^n(t), A)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)} - \frac{\hat{C}^n(t, A)}{\hat{C}^n(t, \mathbb{R}_+)} \right| &= \left| \frac{\hat{C}^n(\Lambda^n(t), A)\hat{C}^n(t, \mathbb{R}_+) - \hat{C}^n(t, A)\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)\hat{C}^n(t, \mathbb{R}_+)} \right| \\ &= \left| \frac{\hat{C}^n(\Lambda^n(t), A)\hat{C}^n(t, \mathbb{R}_+) - \hat{C}^n(t, A)\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)}{\hat{C}^n(\Lambda^n(t), \mathbb{R}_+)\hat{C}^n(t, \mathbb{R}_+)} \right| \\ &\leq \frac{|\hat{C}^n(t, \mathbb{R}_+) - \hat{C}^n(\Lambda^n(t), \mathbb{R}_+)|}{\epsilon} + \frac{|\hat{C}^n(t, A) - \hat{C}^n(\Lambda^n(t), A)|}{\epsilon} \\ &\leq \frac{2\hat{I}^n(t)|t - \Lambda^n(t)|}{\epsilon}, \end{aligned}$$

which implies convergence of both terms to 0 in probability. The convergence of third and fourth terms can be proven similarly, which in turn imply (3.6).

Using the convergence for individual $A \in \mathcal{B}(\mathbb{R}_+)$, we now prove the convergence of the measure-valued processes. To prove our result, we use part 2 of Theorem 5.3 in Mitoma [20], which translated into our case, states that if

1. For each $f \in \mathbb{C}_b(\mathbb{R}_+)$, $\{\langle f, \mathcal{S}_\epsilon^n \eta^n \rangle_t, t \geq 0\}$ is tight in $\mathbb{D}_{\mathbb{R}}[0, T]$.
2. For each $f_1, f_2, \dots, f_m \in \mathbb{C}_b(\mathbb{R}_+)$ and $t_1, t_2, \dots, t_m \in [0, T]$,

$$(\langle f_1, \mathcal{S}_\epsilon^n \eta_{t_1}^n \rangle, \langle f_2, \mathcal{S}_\epsilon^n \eta_{t_2}^n \rangle, \dots, \langle f_m, \mathcal{S}_\epsilon^n \eta_{t_m}^n \rangle) \Rightarrow (\langle f_1, \mathcal{S}_\epsilon \eta_{t_1} \rangle, \langle f_2, \mathcal{S}_\epsilon \eta_{t_2} \rangle, \dots, \langle f_m, \mathcal{S}_\epsilon \eta_{t_m} \rangle) \in \mathbb{R}^m$$

then $\mathcal{S}_\epsilon^n \eta^n \Rightarrow \mathcal{S}_\epsilon \eta$.

We proved the tightness condition in Theorem 2. We prove the second part by proving convergence in probability which is slightly stronger than what is deserved. For any $\varrho > 0$

$$\begin{aligned} \lim_{n \rightarrow \infty} \mathbb{P}(|(\langle f_1, \mathcal{S}_\epsilon^n \eta_{t_1}^n \rangle, \dots, \langle f_m, \mathcal{S}_\epsilon^n \eta_{t_m}^n \rangle) - (\langle f_1, \mathcal{S}_\epsilon \eta_{t_1} \rangle, \dots, \langle f_m, \mathcal{S}_\epsilon \eta_{t_m} \rangle)| > \varrho) \\ \leq \lim_{n \rightarrow \infty} \sum_{i=1}^m \mathbb{P}(|\langle f_i, \mathcal{S}_\epsilon^n \eta_{t_i}^n \rangle - \langle f_i, \mathcal{S}_\epsilon \eta_{t_i} \rangle| > \frac{\varrho}{m}) \end{aligned}$$

The functions $f_i(x) \in \mathbb{C}_b$, i.e., $|f_i(x)| < b_i$ for some b_i and $f_i(x)$ can be approximated by a simple function such that

$$\left| f_i(x) - \sum_{j=1}^{k_i} a_j^i \delta_x(A_j^i) \right| < \frac{\epsilon}{6m}, \text{ for } 1 \leq i \leq m.$$

Therefore,

$$\begin{aligned}
\mathbb{P}(|\langle f_i, \mathcal{S}_\epsilon^n \eta_{t_i}^n \rangle - \langle f_i, \mathcal{S}_\epsilon \eta_{t_i} \rangle| > \frac{\varrho}{m}) &\leq \mathbb{P}(|\langle f_i, \mathcal{S}_\epsilon^n \eta_{t_i}^n \rangle - \langle f_i, \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \rangle| \\
&\quad + |\langle f_i, \mathcal{S}_\epsilon \eta_{t_i} \rangle - \langle f_i, \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \rangle| > \frac{\varrho}{m}) \\
&\leq \mathbb{P}(|\langle f_i, \mathcal{S}_\epsilon^n \eta_{t_i}^n \rangle - \langle f_i, \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \rangle| > \frac{\varrho}{2m}) + \mathbb{P}(2b_i |t_i - \Lambda^n(t_i)| > \frac{\varrho}{m}) \\
&\leq \mathbb{P}\left(\left|\langle f_i, \mathcal{S}_\epsilon^n \eta_{t_i}^n \rangle - \left\langle \sum_{j=1}^{k_i} a_j^i \delta_x(A_j^i), \mathcal{S}_\epsilon^n \eta_{t_i}^n \right\rangle\right| > \frac{\varrho}{6m}\right) \\
&\quad + \mathbb{P}\left(\left|\langle f_i, \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \rangle - \left\langle \sum_{j=1}^{k_i} a_j^i \delta_x(A_j^i), \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \right\rangle\right| > \frac{\varrho}{6m}\right) \\
&\quad + \mathbb{P}\left(\left|\left\langle \sum_{j=1}^{k_i} a_j^i \delta_x(A_j^i), \mathcal{S}_\epsilon^n \eta_{t_i}^n \right\rangle - \left\langle \sum_{j=1}^{k_i} a_j^i \delta_x(A_j^i), \mathcal{S}_\epsilon \eta_{\Lambda^n(t_i)} \right\rangle\right| > \frac{\varrho}{6m}\right) \\
&\quad + \mathbb{P}(2b_i |t_i - \Lambda^n(t_i)| > \frac{\varrho}{m}) \\
&\leq \sum_{j=1}^{k_i} \mathbb{P}\left(|\eta^n(t_i, A_j^i) - \eta(\Lambda^n(t_i), A_j^i)| > \frac{\varrho}{6b_i m k_i}\right) \\
&\quad + \mathbb{P}(2b_i |t_i - \Lambda^n(t_i)| > \frac{\varrho}{m}),
\end{aligned}$$

which converges to 0 as implied by (3.6). \square

Limiting Fairness Measures for Fastest-Server First and Slowest-Server-First Policies

We use Lemma 3 to derive the limiting fairness processes for fastest-server-first and slowest-server-first policies, where the system controller has the knowledge of individual service rates and routes the arriving customer to the fastest and slowest idle servers, respectively.

Theorem 3. *If $\zeta = \delta_{\mu_{\min}}$. where $\mu_{\min} := \text{ess inf}(\tilde{\mu}_k^n)$, then the limiting fairness measure for Fastest-Server-First policy is*

$$\eta_t(A) = \delta_{\mu_{\min}}(A) \text{ for all } t \geq 0.$$

Proof. We prove that for any $\epsilon > 0$ $A^\epsilon = [\mu_{\min} + \epsilon, \infty)$ and $T > 0$

$$\left| \hat{C}^n(t, A^\epsilon) \right|_{*,T} = \hat{C}^n(T, A^\epsilon) \xrightarrow{p} 0.$$

We define L^n as in Lemma 1 and use the fact that it can be partitioned into disjoint intervals as

$$\hat{C}^n(T, A^\epsilon) = \sum_{i=0}^{\infty} \int_{\alpha_i^n}^{\beta_i^n} \frac{\sum_{k=1}^n \delta_{\mu_k}(A^\epsilon) I_k^n(s) ds}{\sqrt{n}} \leq \sum_{i=1}^{\infty} (\beta_i^n - \alpha_i^n) |\hat{I}^n|_{*,t}.$$

Since, we know that $|\hat{I}^n|_{*,t}$ is tight, if we can show that $\sum_{i=1}^{\infty} (\beta_i^n - \alpha_i^n) \xrightarrow{p} 0$, our result will follow. The FSF policy guarantees that if there is an arrival when at least one of the servers with service

rate in A^ϵ is idle, then the new arrival is routed to one of the servers in this set. Hence, within the interval (α_i^n, β_i^n) with $i = 1, 2, \dots$, the number of service completions by the servers with service rate in A^ϵ should be greater than the number of arrivals. Hence,

$$\begin{aligned} \mathbb{P}\left(\sum_{i=1}^{\infty} |\beta_i^n - \alpha_i^n| > 0\right) &\leq \mathbb{P}\left(\sup_{0 < s_1 < s_2 < t} \sum_{k=1}^n \delta_{\mu_k}(A^\epsilon) \left(\frac{S_{P,k}^n(s_2) - S_{P,k}^n(s_1)}{\sqrt{n}}\right) - \frac{A(s_2) - A(s_1)}{\sqrt{n}} > 0\right) \\ &= \mathbb{P}\left(\sup_{0 < s_1 < s_2 < t} \sum_{k=1}^n \delta_{\mu_k}(A^\epsilon) \frac{S_{P,k}^n(s_2) - S_{P,k}^n(s_1) - \tilde{\mu}_k(s_2 - s_1)}{\sqrt{n}} \right. \\ &\quad \left. + \frac{(\sum_{k=1}^n \tilde{\mu}_k - n\bar{\mu})(s_2 - s_1)}{\sqrt{n}} - \frac{\sum_{k=1}^n (1 - \delta_{\tilde{\mu}_k}(A^\epsilon))\tilde{\mu}_k(s_2 - s_1)}{\sqrt{n}} \right. \\ &\quad \left. - \frac{A(s_2) - A(s_1) - \lambda_n(s_2 - s_1)}{\sqrt{n}} + \frac{(n\bar{\mu} - \lambda_n)(s_2 - s_1)}{\sqrt{n}} > 0\right). \end{aligned}$$

The term $\frac{\sum_{k=1}^n (1 - \delta_{\tilde{\mu}_k}(A^\epsilon))\tilde{\mu}_k(s_2 - s_1)}{\sqrt{n}}$ diverges to infinity as $n \rightarrow \infty$. Using the functional central limit theorems for Poisson and renewal processes and continuous mapping theorem for the supremum, we know that the other terms converge in distribution and hence tight. This implies that the probability approaches 0 as $n \rightarrow \infty$. \square

As will be shown in Theorem 5 in the next section, the diffusion approximation provided in Theorem 2.2 in Atar [2] also implies the same limiting fairness measure for the fastest-server-first policy. Our result is slightly stronger as we do not require the support of the service rates to be bounded. However, for Assumption 3 to be satisfied for the slowest-server-first policy, we need to impose the boundedness condition. The rest of the proof of the following theorem is similar to the proof of Theorem 3 and we omit it here.

Theorem 4. *Suppose that $\zeta = \delta_{\mu_{\max}}$, where $\mu_{\max} := \text{ess sup}(\tilde{\mu}_k^n) < \infty$. Then, the limiting fairness measure for Slowest-Server-First policy is*

$$\eta_t(A) = \delta_{\mu_{\max}}(A) \text{ for all } t \geq 0.$$

Remark 1. *Unfortunately, the sequences of fairness processes do not converge in any of the four topologies provided by Skorokhod [25] in general. We try to provide some insight into this claim by considering fastest-server-first policy as an example. Assume $\hat{X}^n(0)$ and their limit ξ_0 is bounded away from 0, hence $\tau_0^n > 0$. We take $\zeta = \delta_{\mu_{\min}}$ and as we prove in Theorem 3 the limiting fairness process is continuous and is also equal to $\delta_{\mu_{\min}}$. Choosing $A = [\mu_{\min} + \epsilon, \infty)$ such that $\mathbb{P}(\tilde{\mu}_k^n \in A) > 0$,*

$$\mathbb{P}\left(\sup_{0 < t < T} |\eta^n(t) - \eta(t)| > \epsilon\right) \geq \mathbb{P}(\exists t > 0 \text{ such that } \eta^n(t) = \delta_{\mu^*} \text{ where } \mu^* \in A | \mathcal{F}_0) = \frac{\sum_{k=1}^n \delta_{\tilde{\mu}_k^n}(A) \tilde{\mu}_k^n}{\sum_{k=1}^n \tilde{\mu}_k^n}.$$

The equality follows using the basic properties of minimum of exponential random variables and realizing that a time t as stated above exists if the first server to become idle is in set A . Hence, it is not possible to find any $\Lambda(\cdot)$ to show convergence of fairness processes in J_1 and J_2 topologies. The distance between graphs of sample paths also is greater than ϵ in the case above, making convergence in M_1 and M_2 topologies impossible.

3.3 Convergence of the System Length Process

In this section, we analyze the weak limit of the scaled system length processes $\hat{X}^n(t)$. Theorem 2 and the completeness of the space of probability measures guarantee that for any sequence of fairness processes, there exists a subsequence with a limiting fairness process. The uniqueness of the limiting process can be proven if the finite dimensional distributions of the process converge. Henceforth, we assume that the sequence of fairness processes has a unique limiting process in the sense in Definition 1.

A general approach in proving limit theorems for queueing systems is to use the structure of Equation (3.2) in conjunction with the continuous mapping theorem (c.f. Theorem 2.7 in Billingsley [6] and Section 4.1 in [22]). The nonstandard definition of the limiting fairness processes in Definition 1 makes it harder for us to use this theorem directly. The proof of the continuous mapping theorem relies on the Skorokhod Representation Theorem for which we provide an extension in the following lemma.

Lemma 4. *Let S_k be Polish spaces and $\Xi^n = (X_1^n, X_2^n, \dots)$ be random elements where $X_k^n \in S_k$ for all $k, n \in \mathbb{N}$ and $\{X_{k,n}\}_{n \in \mathbb{N}}$ is tight for each k . There exist $n_1 < n_2 < \dots$ in \mathbb{N} and random variables Y_{n_j} defined on the same probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$, taking values in S_{n_j} such that for each j the family $(Y_{k,j})_{k \in \mathbb{N}}$ has the same joint law with $(X_{k,n_j})_{k \in \mathbb{N}}$ and for each $k \in \mathbb{N}$, $Y_{k,j}$ converges almost surely as $j \rightarrow \infty$.*

Proof. Let $\tilde{S} = S_1 \times S_2 \times \dots$, an infinite Cartesian product, which is a Polish space. Then $\tilde{X}_n = (X_{1n}, X_{2n}, \dots)$ are \tilde{S} -valued random elements for each $n \in \mathbb{N}$. As $\{X_k^n\}_{n \in \mathbb{N}}$ is tight, there exists a compact set $E_k \subset S_k$ such that $\mathbb{P}(X_{kn} \notin E_k \text{ for some } n) < 2^{-k}\epsilon$ for each $k \in \mathbb{N}$. This implies that $\mathbb{P}(X_{kn} \notin E_k \text{ for some } k, n \in \mathbb{N}) < \epsilon$. Using Tychonoff's theorem (see Chapter 5, Munkres [21]), $E_1 \times E_2 \times \dots$ is compact in \tilde{S} . Hence, we conclude that $\{\tilde{X}_n\}_{n \in \mathbb{N}}$ is also tight. Hence, using Skorokhod's representation theorem, we can find $n_1 < n_2 < \dots$ and \tilde{Y}_j on a probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$ which has the same distribution as \tilde{X}_{n_j} and converges almost surely as $j \rightarrow \infty$. Hence, the lemma follows. \square

Now, we are ready to prove our main result.

Theorem 5. *Suppose that $\{\eta_t\}_{t \in \mathbb{R}}$ is the limiting fairness process for the sequence of queueing systems as defined in Definition 1. Then, the scaled process \hat{X}^n converges weakly to ξ , where ξ is the solution of*

$$\xi(t) = \xi_0 + (\bar{\mu}\sqrt{c_a^2 + 1})W(t) + \beta t + \langle \iota, \eta_t \rangle \int_0^t \xi^-(s)ds - \gamma \int_0^t \xi^+(s)ds, \text{ for all } t \geq 0, \quad (3.7)$$

where $c_a^2 := \text{var}(u_k^n)/\bar{\mu}^2$, i.e., the coefficient of variation of u_k^n s, $\beta := \hat{\lambda} - \tilde{\sigma}$, $\tilde{\sigma}$ is distributed normally with mean 0 and variance $\text{var}(\tilde{\mu})$.

Proof. We focus on Equation (3.2) and use the martingale method outlined in [22]. The processes

$$\begin{aligned} \hat{M}_1^n(t) &:= \frac{\sum_{k=1}^n D_k^n(t) - \tilde{\mu}_k^n \int_0^t (1 - I_k^n(s))ds}{\sqrt{n}} \\ \hat{M}_2^n(t) &:= \frac{N^n \left(\gamma \int_0^t (X^n(s) - n)^+ ds \right) - \gamma \int_0^t (X^n(s) - n)^+ ds}{\sqrt{n}}, \end{aligned}$$

corresponding to the fourth and sixth terms in (3.2) are martingales with predictable quadratic variations

$$\begin{aligned}\langle \hat{M}_1^n \rangle_t &= \frac{\sum_{k=1}^n \tilde{\mu}_k^n \int_0^t (1 - I_k^n(s)) ds}{n} \\ \langle \hat{M}_2^n \rangle_t &= \frac{\gamma \int_0^t (X^n(s) - n)^+ ds}{n},\end{aligned}$$

respectively. Using Corollary 1,

$$\frac{\sum_{k=1}^n \tilde{\mu}_k^n t - \sum_{k=1}^n \tilde{\mu}_k^n \int_0^t I_k^n(s) ds}{n} \xrightarrow{p} \bar{\mu} t \text{ in } D_{\mathbb{R}_+}[0, \infty) \text{ and } \left| \frac{\int_0^t (X^n(s) - n)^+ ds}{n} \right|_{*,T} \xrightarrow{p} 0.$$

Hence, using the martingale central limit theorem

$$\hat{M}_1^n(t) \Rightarrow \bar{\mu} W_2(t) \text{ and } \hat{M}_2^n(t) \xrightarrow{p} 0.$$

And as a result of central limit theorem for renewal processes and standard central limit theorem, we have

$$\hat{A}^n(t) = \frac{A^n(t) - \lambda_n t}{\sqrt{n}} \Rightarrow c_a \bar{\mu} W_1(t) \text{ and } \hat{\sigma}^n = \frac{\tilde{\mu}_k^n - \bar{\mu}}{\sqrt{n}} \Rightarrow \tilde{\sigma}.$$

Taking $\epsilon_k = 1/k$ for all $k \in \mathbb{N}$ and $\Xi^n = (\hat{X}^n(0), \hat{A}^n, \hat{M}_1^n, \hat{M}_2^n, \hat{\sigma}^n, \mathcal{S}_{\epsilon_1}^n \eta^n, \mathcal{S}_{\epsilon_2}^n \eta^n, \dots)$, Lemma 4 implies that we can find $\check{\Xi}^k = (\check{X}^k(0), \check{A}^k, \check{M}_1^k, \check{M}_2^k, \check{\sigma}^k, \mathcal{S}_{\epsilon_1}^k \check{\eta}^k, \mathcal{S}_{\epsilon_2}^k \check{\eta}^k, \dots)$ such that $\check{\Xi}^k$ has the same distribution as Ξ^{n_k} and

$$\check{\Xi}^k \rightarrow (\check{\xi}_0, c_a \bar{\mu} \check{W}_1, \bar{\mu} \check{W}_2, 0, \check{\sigma}, \mathcal{S}_{\epsilon_1} \check{\eta}, \mathcal{S}_{\epsilon_2} \check{\eta}, \dots),$$

where $(\check{\xi}_0, c_a \bar{\mu} \check{W}_1, \bar{\mu} \check{W}_2, 0, \check{\sigma}, \mathcal{S}_{\epsilon_1} \check{\eta}, \mathcal{S}_{\epsilon_2} \check{\eta}, \dots)$ has the same joint distribution as

$$(\xi_0, c_a \bar{\mu} W_1, \bar{\mu} W_2, 0, \tilde{\sigma}, \mathcal{S}_{\epsilon_1} \eta, \mathcal{S}_{\epsilon_2} \eta, \dots).$$

As the convergence occurs in Skorokhod- J_1 topology, Theorem 4.1 in [28] implies that for any fixed $\epsilon > 0$ we can find a common sequence of homeomorphisms $\Lambda^n : [0, T] \rightarrow [0, T]$ such that with probability 1,

$$\left| \check{A}^n(t) - \bar{\mu} c_a \check{W}_1(\Lambda^n(t)) \right|_{*,T} \vee \left| \check{M}_1^n(t) - \bar{\mu} \check{W}_2(\Lambda^n(t)) \right|_{*,T} \vee \left| \langle \iota, \mathcal{S}_{\epsilon} \check{\eta}_t^n \rangle - \langle \iota, \mathcal{S}_{\epsilon} \check{\eta}_{\Lambda^n(t)} \rangle \right|_{*,T} \vee \left| \dot{\Lambda}^n(t) - 1 \right|_{*,T} \rightarrow 0,$$

as $n \rightarrow \infty$, where $\dot{\Lambda}^n$ denotes the derivative of Λ with respect to t . Also, Assumption 3 and the tightness proved in Lemma 1 imply that for any $\rho > 0$, there exists an $K_\rho > 0$ such that

$$\mathbb{P} \left(\sup_{n \in \mathbb{N}} \left\{ \left| \langle \iota, \check{\eta}_t^n \rangle \right|_{*,T} \vee \left| \check{X}^n(t) \right|_{*,T} \right\} > K_\rho \right) < \rho.$$

Let $\check{\xi}(t)$ be the unique strong solution of

$$\check{\xi}(t) = \check{\xi}_0 + c_a \bar{\mu} \check{W}_1(t) + \bar{\mu} \check{W}_2(t) + \beta t + \langle \iota, \check{\eta}_t \rangle \int_0^t \check{\xi}^-(s) ds - \gamma \int_0^t \check{\xi}^+(s) ds, \text{ for all } t \geq 0.$$

Defining \check{X} by replacing Ξ with $\check{\Xi}$ in (3.2), we need to prove that for any $\rho, \varrho > 0$, we can find a $N_{\rho, \varrho}$ such that $n > N_{\rho, \varrho}$ implies

$$\mathbb{P}(d_S(\check{X}^n, \check{\xi}) > \varrho) < \rho.$$

Choosing $K_{\rho/2}$ as defined above

$$\mathbb{P}\left(d_S(\check{X}^n, \check{\xi}) > \varrho\right) \leq \mathbb{P}\left(d_S(\check{X}^n, \check{\xi}) > \varrho, \sup_{n \in \mathbb{N}} \left\{ \|\langle \iota, \check{\eta}_t^n \rangle\|_{*,T} \vee \left| \check{X}^n(t) \right|_{*,T} \right\} \leq K_{\rho/2}\right) + \frac{\rho}{2}. \quad (3.8)$$

Hence, we concentrate on the scenarios corresponding to the first term on the righthand side of (3.8) and assume

$$\sup_{n \in \mathbb{N}} \left\{ \|\langle \iota, \check{\eta}_t^n \rangle\|_{*,T} \vee \left| \check{X}^n(t) \right|_{*,T} \right\} \leq K_{\rho/2}$$

For any $\varepsilon > 0$, there is a sufficiently large N_0 , such that for any $n > N_0$

$$\begin{aligned} |\check{X}^n(t) - \check{\xi}(\Lambda^n(t))| &\leq \varepsilon + \gamma \left| \int_0^t (\check{X}^n(s))^+ ds - \int_0^{\Lambda^n(t)} \check{\xi}(s)^+ ds \right| \\ &\quad + \left| \langle \iota, \check{\eta}_t^n \rangle \int_0^t (\check{X}^n(s))^- ds - \langle \iota, \check{\eta}_{\Lambda^n(t)} \rangle \int_0^{\Lambda^n(t)} \check{\xi}(s)^- ds \right| \\ &\leq \varepsilon + \gamma \left| \int_0^t \left((\check{X}^n(s))^+ - \check{\xi}(\Lambda^n(s))^+ \right) ds + \int_0^t (1 - \dot{\Lambda}^n(s)) \check{\xi}(\Lambda^n(s))^+ ds \right| \\ &\quad + \left| (\langle \iota, \check{\eta}_t^n \rangle - \langle \iota, \mathcal{S}_\epsilon \check{\eta}_t^n \rangle) \int_0^t (\check{X}^n(s))^- ds \right|_t \\ &\quad \quad - (\langle \iota, \check{\eta}_{\Lambda^n(t)} \rangle - \langle \iota, \mathcal{S}_\epsilon \check{\eta}_{\Lambda^n(t)} \rangle) \int_0^{\Lambda^n(t)} \check{\xi}(s)^- ds \\ &\quad + \left| (\langle \iota, \mathcal{S}_\epsilon \check{\eta}_t^n \rangle - \langle \iota, \mathcal{S}_\epsilon \check{\eta}_{\Lambda^n(t)} \rangle) \int_0^t (\check{X}^n(s))^- ds \right| \\ &\quad \quad - \langle \iota, \mathcal{S}_\epsilon \check{\eta}_{\Lambda^n(t)} \rangle \int_0^t \left((\check{X}^n(s))^- - \check{\xi}(\Lambda^n(s))^- \right) ds \\ &\quad + \left| \langle \iota, \mathcal{S}_\epsilon \check{\eta}_{\Lambda^n(t)} \rangle \int_0^t \dot{\Lambda}^n(s) \check{\xi}(\Lambda^n(s))^- ds \right| \\ &\leq \varepsilon + \gamma \int_0^t \left| (\check{X}^n(s)) - \check{\xi}(\Lambda^n(s)) \right| ds + \varepsilon K_{\rho/2} T + 4\epsilon K_{\rho/2} + \varepsilon K_{\rho/2} \\ &\quad + K_{\rho/2} \int_0^t \left| \check{X}^n(s) - \check{\xi}(\Lambda^n(s)) \right| ds + \varepsilon (K_{\rho/2})^2 T \end{aligned}$$

Now, using Gronwall's inequality (c.f. Lemma 4.1 in [22]), we have

$$d_S(\check{X}^n, \check{\xi}) \leq \sup_{0 \leq t \leq T} |\check{X}^n(t) - \check{\xi}(\Lambda^n(t))| \leq (\varepsilon(1 + K_{\rho/2}(1 + T + K_{\rho/2}T)) + 4\epsilon K_{\rho/2}) e^{(\gamma + K_{\rho/2})T}$$

Hence, choosing ε and ϵ appropriately and N_0 large enough, we prove our result. \square

4 Totally Blind Routing Policies

4.1 A Related Martingale

In this section, our goal is to identify fairness processes corresponding to some specific policies. The key term in defining the fairness process is the cumulative idle time by time t experienced by servers whose service rate fall in set A . To analyze the total idleness process further, we use the doubly stochastic process $\{S_P^n(t)\}_{t \geq 0}$ that represents the potential service completion times and define

$$\phi_i^n = \inf\{t - \theta_i^n : I_{\kappa_i^n}(t) = 0, t > \theta_i^n\} \text{ for all } i \in \mathbb{N} \text{ and } n \in \mathbb{N}.$$

The random variable ϕ_i^n represents the length of the idle period experienced by the server who becomes idle at time θ_i^n . If an actual service completion occurs at θ_i^n , then $I_{\kappa_i^n}(\theta_i^n-) = 1$ and $\phi_i^n > 0$, and if θ_i^n is not an actual service completion time, i.e., $I_{\kappa_i^n}(\theta_i^n-) = 0$, then $\phi_i^n = 0$. Similarly, we define

$$\phi_{-k}^n = \inf\{t > 0 : I_k^n(t) = 0\}, \text{ for all } 1 \leq k \leq n \text{ and } n \in \mathbb{N},$$

which represents the duration of the idle period experienced by servers who are idle at time 0. Hence, defining

$$\Upsilon_i^n(A) = \begin{cases} 1 & \text{if } \tilde{\mu}_{\kappa_i^n} \in A \\ 0 & \text{otherwise,} \end{cases}$$

we can equivalently represent the total cumulative idleness experienced by servers in A by time t as

$$\sum_{k=1}^n \delta_{\tilde{\mu}_k^n}(A) \int_0^t I_k^n(s) ds = \sum_{k=1}^n \delta_{\tilde{\mu}_k^n}(A) (\phi_{-k}^n \wedge t) + \sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n \wedge (t - \theta_i^n)).$$

To analyze the righthand side further, we introduce the following martingale:

Lemma 5. *For any Borel set $A \subset \mathbb{R}_+$,*

$$\begin{aligned} M_A^n(t) = n^{-1/2} & \left(\sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n \wedge (t - \theta_i^n)) - \sum_{i=1}^{S_P^n(t)} \mathbb{E} \left[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n \right] \right. \\ & \left. + \sum_{i=1}^{S_P^n(t)} (\Upsilon_i^n(A) \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t^n]) \right) \end{aligned} \quad (4.1)$$

is an \mathbf{F}^n -martingale.

Proof. Being conditional expectations, second and third sums on the right-hand side of (4.1) are clearly \mathcal{F}_t -measurable. The random variables $(\phi_i^n + \theta_i^n)$ are stopping times for any $0 < i \leq S_P(t)$, which implies that the first sum is also measurable with respect to \mathcal{F}_t . For any $n \geq 1$ and $t > \theta_i$, we have

$$0 \leq X^n(t) < n \text{ and } \phi_i^n - \theta_i^n < \inf\{s > t : X^n(s) = n\}.$$

Considering the discrete-time Markov chain $Y_k^n = X^n(\sum_{i=1}^k u_i^n / \lambda^n)$ and the positive recurrence of this chain when restricted to be between 0 and n proves that $\mathbb{E}[|M_t(A)^n|] < \infty$ for all $t \geq 0$. The

rest of the proof relies on the basic relationship that $\phi_i^n \wedge (t - \theta_i^n) = \phi_i^n - (\phi_i^n + \theta_i^n - t)^+$. We have

$$\begin{aligned}
\mathbb{E}[M_A^n(t+s) - M_A^n(t) | \mathcal{F}_t] &= \mathbb{E} \left[\sum_{i=1}^{S_P^n(t+s)} \Upsilon_i^n(A) \phi_i^n - \sum_{i=1}^{S_P^n(t+s)} \Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ \right. \\
&\quad - \sum_{i=1}^{S_P^n(t+s)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}] + \sum_{i=1}^{S_P^n(t+s)} \mathbb{E}[\Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ | \mathcal{F}_{t+s}] \\
&\quad - \sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) \phi_i^n + \sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n - t + \theta_i^n)^+ \\
&\quad \left. + \sum_{i=1}^{S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}] - \sum_{i=1}^{S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) (\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t] \middle| \mathcal{F}_t \right] \\
&= \mathbb{E} \left[\sum_{i=S_P^n(t)+1}^{S_P^n(t+s)} \Upsilon_i^n(A) \phi_i^n - \sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ \right. \\
&\quad - \sum_{i=S_P^n(t)+1}^{S_P^n(t+s)} \Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ - \sum_{i=S_P^n(t)+1}^{S_P^n(t+s)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}] \\
&\quad + \sum_{i=1}^{S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ | \mathcal{F}_{t+s}] \\
&\quad + \sum_{i=S_P^n(t)+1}^{S_P^n(t+s)} \mathbb{E}[\Upsilon_i^n(A) (\phi_i^n - t - s + \theta_i^n)^+ | \mathcal{F}_{t+s}] + \sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n - t + \theta_i^n)^+ \\
&\quad \left. - \sum_{i=1}^{N^n(t)} \mathbb{E}[\Upsilon_i^n(A) (\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t] \middle| \mathcal{F}_t \right] = 0
\end{aligned}$$

This proves the lemma. \square

The $\hat{C}^n(t, A)$ is a continuous function of time t for any $A \in \mathcal{B}(\mathbb{R}_+)$ and hence a previsible process. Hence, in Lemma 5, rather than finding a compensator for the cumulative total idleness process, we find a process who is compensator is the cumulative total idleness process and hence there is no uniqueness claim.

Lemma 6. *If for any $T > 0$, there exist random variables $\vartheta_{n,T}$ such that $\mathbb{E}[\vartheta_T^n] \rightarrow 0$ as $n \rightarrow \infty$ and*

$$\sup_{0 \leq t \leq T} \sup_{-n \leq i \leq S_P^n(t)} \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t^n] < \vartheta_T^n, \quad (4.2)$$

then

$$\sup_{0 \leq t \leq S_P^n(T)} \left| \hat{C}^n(t, A) - \frac{\sum_{i=1}^{S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} \right| \xrightarrow{P} 0.$$

Proof. To prove our result we show that the martingale M_A^n and the third term on the righthand side of (4.1) converge to 0 in probability under the stated conditions. Condition (4.2) imply that for all $T \geq 0$ and $n \in \mathbb{N}$, ϑ_T^n is positive and converges 0 in probability using Markov inequality. If the expected maximum jump and the optional quadratic variation of M_A^n converges 0, the latter being in probability, the martingale can be shown to converges 0 in probability using the martingale central limit theorem (c.f. Theorem 8.1 in [22]). We first note that (4.2) also implies

$$\sup_{1 \leq i \leq S_P^n(T)} \mathbb{E} \left[\phi_i^n | \mathcal{F}_{\theta_i^n-}^n \right] \xrightarrow{p} 0,$$

which also imply the expected maximum jump of M_A^n converges 0. Expanding the optional quadratic variation of M_A^n , we get

$$\begin{aligned} [M_A^n] \leq & \frac{2 \left[\sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A) (\phi_i^n \wedge (t - \theta_i^n)) \right]}{n} + \frac{2 \left[\sum_{i=1}^{S_P^n(t)} \mathbb{E} \left[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n \right] \right]}{n} \\ & + \frac{2 \left[\sum_{i=1}^{S_P^n(t)} (\Upsilon_i^n(A) \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t^n]) \right]}{n}. \end{aligned} \quad (4.3)$$

The first term on the righthand side of (4.3) is the optional quadratic variation of a continuously increasing process and hence is 0 for all $T \geq 0$ and $n \in \mathbb{N}$. The second term is the optional quadratic variation of an increasing pure jump process and

$$\frac{\left[\sum_{i=1}^{S_P^n(t)} \mathbb{E} \left[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n \right] \right]}{n} = \frac{\sum_{i=1}^{S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n]^2}{n} \leq \left(\sup_{1 \leq i \leq S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n] \right)^2 \frac{S_P^n(t)}{n} \xrightarrow{p} 0.$$

For analyzing the third term, we define $\varpi_i^n(t) = \inf\{I_{\kappa_i}(s) : \theta_i^n \leq s \leq t\}$, for $i \leq S_P^n(t)$, $n \in \mathbb{N}$, which is 1 if the server completing the service at θ_i^n stays idle until time t and is 0 otherwise. The function $\varpi_i^n(t)$ enables us to explicitly address only the servers who are idle at time t as for any $\theta_i^n \leq S_P^n(t)$,

$$\mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t] = \varpi_i(t) \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t].$$

We now analyze the optional quadratic variation of the third term on the righthand side of (4.3) directly using the definition (c.f. Theorem 3.3 in [22]). Consider any finite partition of $[0, T]$ such

that $0 = t_0 < t_1 < \dots < t_m = T$ and

$$\begin{aligned}
& n^{-1} \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_l)} \Upsilon_i(A) \varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right)^2 \\
&= n^{-1} \sum_{l=1}^m \left(\sum_{i=S_P(t_{l-1}+1)}^{S_P(t_l)} \Upsilon_i(A) \varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] \right. \\
&\quad \left. + \sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (\varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]) \right)^2 \\
&\leq 2n^{-1} \sum_{l=1}^m \left(\sum_{i=S_P(t_{l-1}+1)}^{S_P(t_l)} \Upsilon_i(A) \varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] \right)^2 \\
&\quad + 2n^{-1} \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (\varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]) \right)^2
\end{aligned}$$

To analyze the second term further, we realize that $\varpi_i(t_l) \leq \varpi_i(t_{l-1})$ for any $i < S_P(t_{l-1})$ and write

$$\begin{aligned}
& \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (\varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]) \right)^2 \\
&= \sum_{l=1}^m \left(- \sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (1 - \varpi_i(t_l)) \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right. \\
&\quad \left. + \sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (\varpi_i(t_l)) \varpi_i(t_{l-1}) (\mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]) \right)^2 \\
&\leq 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (1 - \varpi_i(t_l)) \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right)^2 \\
&\quad + 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (\varpi_i(t_l)) \varpi_i(t_{l-1}) (\mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]) \right)^2
\end{aligned}$$

Again considering the second term on the righthand side, we have

$$\begin{aligned}
& \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A)(\varpi_i(t_l))\varpi_i(t_{l-1}) \left(\mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right) \right)^2 \\
& \leq 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A)(\varpi_i(t_l))\varpi_i(t_{l-1}) \left(\mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_l}] \right) \right)^2 \\
& \quad + 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A)(\varpi_i(t_l))\varpi_i(t_{l-1}) \left(\mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right) \right)^2 \\
& \leq 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A)(\varpi_i(t_l))\varpi_i(t_{l-1}) (t_l - t_{l-1}) \right)^2 \\
& \quad + 2 \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A)(\varpi_i(t_l))\varpi_i(t_{l-1}) \left(\mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right) \right)^2
\end{aligned}$$

Aggregating all our calculations above, we get

$$\frac{\left[\sum_{i=1}^{S_P^n(t)} (\Upsilon_i^n(A) \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t^n]) \right]}{n} \leq 2V_1^n(t) + 4V_2^n(t) + 8(V_3^n(t) + V_4^n(t)),$$

where

$$\begin{aligned}
V_1^n(t) &= \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m \left(\sum_{i=S_P(t_{l-1}+1)}^{S_P^n(t_l)} \Upsilon_i(A)\varpi_i(t_l) \mathbb{E}[(\phi_i^n - t_l + \theta_i)^+ | \mathcal{F}_{t_l}] \right)^2}{n}, \\
V_2^n(t) &= \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m \left(\sum_{i=1}^{S_P^n(t_{l-1})} \Upsilon_i(A)(1 - \varpi_i(t_l))\varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right)^2}{n}, \\
V_3^n(t) &= \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m \left(\sum_{i=1}^{S_P^n(t_{l-1})} \Upsilon_i(A)\varpi_i(t_l)\varpi_i(t_{l-1}) (t_l - t_{l-1}) \right)^2}{n}, \\
V_4^n(t) &= \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m \left(\sum_{i=1}^{S_P^n(t_{l-1})} \Upsilon_i(A)\varpi_i(t_l)\varpi_i(t_{l-1}) \left(\mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_l}] - \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}] \right) \right)^2}{n}.
\end{aligned}$$

The first term V_1^n is the definition of the quadratic variation of a pure jump process and hence,

$$V_1^n(t) = \frac{\sum_{i=1}^{S_P(t)} (\Upsilon_i(A) \mathbb{E}[\phi_i^n | \mathcal{F}_{\theta_i^n}])^2}{n} = \left(\sup_{1 \leq i \leq S_P^n(t)} \mathbb{E}[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n}] \right)^2 \frac{S_P^n(t)}{n} \xrightarrow{p} 0.$$

For $V_2^n(t)$,

$$V_2^n(t) \leq (\vartheta_T^n)^2 \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m \left(\sum_{i=1}^{S_P(t_{l-1})} \Upsilon_i(A) (1 - \varpi_i(t_l)) \varpi_i(t_{l-1}) \right)^2}{n}$$

The fraction above defines the optional quadratic variation of a counting process, which counts the number of customers routed to a server by time t . This counting process is bounded by the arrival process $A^n(t)$ and hence (4.2) implies that $V_2^n(t) \xrightarrow{P} 0$ in $\mathbb{D}_{\mathbb{R}}[0, T]$.

For $V_3^n(t)$, write

$$\begin{aligned} V_3^n(t) &\leq \frac{\lim_{\sup|t_l - t_{l-1}| \rightarrow 0} S_P(t) \sum_{l=1}^m \sum_{i=1}^{S_P(t_{l-1})} (\Upsilon_i(A) \varpi_i(t_l) \varpi_i(t_{l-1}) (t_l - t_{l-1}))^2}{n} \\ &\leq \frac{2S_P(t) \sum_{i=1}^{S_P(t)} \lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m (t_l - t_{l-1})^2}{n} = 0 \end{aligned}$$

For the analysis of $V_4^n(t)$, we use a similar approach but this time making explicit use of the fact that the non-zero terms in the summation at any t_{l-1} is bounded by the number of idle servers at that time. Also,

$$\begin{aligned} V_4^n(t) &\leq \lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m 2I^n(t_l) \sum_{i=1}^{S_P(t_{l-1})} \left(\frac{\Upsilon_i(A) \varpi_i(t_l) \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_l}]}{n} \right. \\ &\quad \left. - \frac{\Upsilon_i(A) \varpi_i(t_l) \varpi_i(t_{l-1}) \mathbb{E}[(\phi_i^n - t_{l-1} + \theta_i)^+ | \mathcal{F}_{t_{l-1}}]}{n} \right)^2 \\ &\leq \frac{2|I^n|_{t,*}^2 \lim_{\sup|t_l - t_{l-1}| \rightarrow 0} \sum_{l=1}^m (\mathbb{E}[(\phi_i^n | \mathcal{F}_{t_l}] - \mathbb{E}[\phi_i^n | \mathcal{F}_{t_{l-1}}])^2}{n} \end{aligned}$$

Theorem 2 imply that $|I^n|_{t,*}^2/n$ is tight. The limit and summation is the optional quadratic variation of the square integrable martingale $E[\phi_i^n | \mathcal{F}_t]$ and using Burkholder-Davis-Gundy inequality (Theorem iV.48 in [23]) with $p = 1$

$$\mathbb{E}[\phi_i^n | \mathcal{F}_t^n] \leq c_1 \mathbb{E} \left[\sup_{0 \leq s \leq t} \sup_{1 \leq i \leq S_P^n(t)} \mathbb{E}[\phi_i^n | \mathcal{F}_s^n] \right] \leq c_1 \mathbb{E}[\vartheta_T^n]$$

and hence the optional quadratic variation of $E[\phi_i^n | \mathcal{F}_t]$ converges 0 in probability. This proves that $[M_A^n] \xrightarrow{P} 0$ in probability and using the martingale functional central limit theorem we conclude that $M_A^n(t) \xrightarrow{P} 0$. Using a similar approach, the term on the righthand side of (4.1)

$$\sum_{i=1}^{S_P^n(t)} (\Upsilon_i^n(A) \mathbb{E}[(\phi_i^n - t + \theta_i^n)^+ | \mathcal{F}_t^n]) \leq |\hat{I}^n|_{T,*} \vartheta_T^n \xrightarrow{P} 0.$$

This implies

$$n^{-1/2} \left(\sum_{i=1}^{S_P^n(t)} \Upsilon_i^n(A)(\phi_i^n \wedge (t - \theta_i^n)) - \sum_{i=1}^{S_P^n(t)} \mathbb{E} \left[\Upsilon_i^n(A) \phi_i^n | \mathcal{F}_{\theta_i^n-}^n \right] \right) \xrightarrow{p} 0 \text{ in } \mathbb{D}_{\mathbb{R}}[0, T].$$

Since the limiting function is constant and hence continuous, we conclude that the same convergence also holds with the supremum norm, which proves the lemma. \square

4.2 The Fairness Measure for Totally Blind Policies

We are now ready to address specific policies when service rates are random. Atar [2] proves that for the Longest-Idle-Server-First policy, the system length process converges to a diffusion as given in (1), where

$$\langle \iota, \eta \rangle_t = \frac{\int_0^\infty \mu^2 dF(\mu)}{\int_0^\infty \mu dF(\mu)}, \text{ for all } t \geq 0.$$

This implies that the limiting fairness process is constant through time and

$$\eta_t(A) = \frac{\int_A \mu dF(\mu)}{\int_0^\infty \mu dF(\mu)}, \text{ for all } t \geq 0.$$

In this section, under some mild conditions which readily holds for the LISF policy, we prove that the scaled system length process converges to the same limiting diffusion for a more general class of policies, which we call to be “totally blind policies”. Ward and Armony [26] describe a blind policy to be a control policy that can depend on the system state but not the system parameters. According to this definition and its usage in the literature, a blind policy is assumed not to have perfect information of parameters, however it can make decisions that have implicit dependencies on parameters through the system state. For our purposes we need to define the notion of blindness a bit more restrictively, to mean the expected idling time of a server who becomes idle at time θ_i^n is asymptotically the same with or without the knowledge of the service rate of the server.

Definition 2. *A routing policy is a totally blind policy, if and only if*

$$\sup_{1 \leq i \leq S_P^n(T)} \sqrt{n} \left| \mathbb{E}[\phi_i^n | \mathcal{F}_{\theta_i^n-}^n] - \mathbb{E}[\phi_i^n | \mathcal{F}_{\theta_i^n}^n] \right| \xrightarrow{p} 0$$

Theorem 6. *Under any totally blind policy, i.e.,*

$$\lim_{n \rightarrow \infty} \sup_{1 \leq i \leq S_P^n(T)} \sqrt{n} \left| \mathbb{E}[\phi_i^n | \mathcal{F}_{\theta_i^n-}^n] - \mathbb{E}[\phi_i^n | \mathcal{F}_{\theta_i^n}^n] \right| \xrightarrow{p} 0, \text{ for all } T > 0,$$

where condition (4.2) holds, then

$$\eta_t(A) = \frac{\int_A \mu dF(\mu)}{\int_0^\infty \mu dF(\mu)}, \text{ for all } t \geq 0.$$

Proof. First, we analyze

$$\begin{aligned}
\sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} &= \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\mathbb{E}[\Upsilon_i(A)\phi_i|\mathcal{F}_{\theta_i^n-}]]}{\sqrt{n}} \\
&= \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}]]}{\sqrt{n}} \\
&= \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)(\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} + \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}]]}{\sqrt{n}} \\
&= \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)(\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} + \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)|\mathcal{F}_{\theta_i^n-}]\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}}
\end{aligned}$$

Hence, we have

$$\sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} - \frac{\mathbb{E}[\Upsilon_i(A)|\mathcal{F}_{\theta_i^n-}]\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} = \sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)(\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}}$$

For any $\epsilon > 0$

$$\begin{aligned}
&\mathbb{P}\left(\left|\sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)(\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}}\right| > \epsilon_1\right) \\
&= \mathbb{P}\left(\left|\sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[\Upsilon_i(A)(\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}}\right| > \epsilon_1, \frac{|S_P^n(T) - n\bar{\mu}T|}{n} < \epsilon_2\right) \\
&\quad + \mathbb{P}\left(\frac{|S_P^n(T) - n\bar{\mu}T|}{n} > \epsilon_2\right) \\
&\leq \mathbb{P}\left(\sum_{i=1}^{S_P^n(t)} \frac{\mathbb{E}[|\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])||\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} > \epsilon_1, \frac{|S_P^n(T) - n\bar{\mu}T|}{n} < \epsilon_2\right) \\
&\quad + \mathbb{P}\left(\frac{|S_P^n(T) - n\bar{\mu}T|}{n} > \epsilon_2\right) \\
&\leq \mathbb{P}\left(\frac{(\bar{\mu}T + \epsilon_2)\sqrt{n} \sup_{1 \leq i \leq S_P(T)} \mathbb{E}[|\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}] - \mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}])||\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} > \epsilon_1\right) \\
&\quad + \mathbb{P}\left(\frac{|S_P^n(T) - n\bar{\mu}T|}{n} > \epsilon_2\right)
\end{aligned}$$

This implies

$$\sum_{i=1}^{S_P^n(t)} \left(\frac{\mathbb{E}[\Upsilon_i(A)\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} - \frac{\mathbb{E}[\Upsilon_i(A)|\mathcal{F}_{\theta_i^n-}]\mathbb{E}[\phi_i|\mathcal{F}_{\theta_i^n-}]}{\sqrt{n}} \right) \xrightarrow{p} 0 \text{ in } \mathbb{D}_{\mathbb{R}}[0, T].$$

Again, since the limiting function (zero function) is constant and hence continuous, this convergence can be taken in the supremum norm. Combining this with Lemma 6 and realizing that

$$\mathbb{E}[\Upsilon_i^n(A)|\mathcal{F}_{\theta_i^n-}] = \frac{n^{-1} \sum_{k=1}^n \delta_{\tilde{\mu}_k^n}(A)(1 - I_k^n(\theta_i^n-))\tilde{\mu}_k^n}{n^{-1} \sum_{k=1}^n \tilde{\mu}_k^n} \rightarrow \frac{\int_A \mu dF(\mu)}{\int_0^\infty \mu dF(\mu)} \text{ w.p. } 1,$$

we have

$$\sup_{0 \leq s \leq T} \left| \hat{C}^n(t, A) - \frac{\int_A \mu dF(\mu)}{\int_0^\infty \mu dF(\mu)} \int_0^t \hat{I}(s) ds \right| \xrightarrow{p} 0.$$

Our theorem is then implied by Lemma 2. \square

Intuitively, as the number of servers approaches infinity, the number of idle servers are negligible compared to the servers who are busy. It is well known from basic probability that if we have two independent exponential random variables, say U and V , with rates λ_1 and λ_2 , the probability that U is less than V is $\lambda_1/(\lambda_1 + \lambda_2)$. Hence, using this basic property and the fact that the idle servers are negligible, the probability that a server becoming idle at any given time belongs to set A is asymptotically constant in time and equal to

$$\frac{\int_A \mu dF(\mu)}{\int_0^\infty \mu dF(\mu)}.$$

In other words, totally blind policies asymptotically equalize the time to stay idle for all servers once they become idle. The servers with higher service rates become idle more frequently, proportional to their service rates and hence, their share of total cumulative idle time is also proportional to their service rates.

Longest-Idle-Server-First Policy

We show that Longest-Idle-Server-First policy is totally blind and the condition 4.2 holds. We know that the server who becomes idle at time θ_i^n starts the next service exactly after arrival of the next $\sum_{k=1}^n I_k^n(\theta_i^n)$ customers. Hence,

$$\mathbb{E}[\phi_i^n|\mathcal{F}_t] \leq n^{-1} \left(\sup_{0 \leq s \leq T} \mathbb{E}[u_{A^n(s)+1}^n|\mathcal{F}_s] + \mathbb{E}[u_1^n] \sup_{0 \leq s \leq T} \sum_{k=1}^n I_k^n(s) \right).$$

and

$$\mathbb{E}[\phi_i^n|\mathcal{F}_{\theta_i^n}] = \mathbb{E}[\phi_i^n|\mathcal{F}_{\theta_i^n-}] \quad (4.4)$$

for all $i, n \in \mathbb{N}$, as it only depends on the number of idle servers at θ_i^n , which is $\mathcal{F}_{\theta_i^n-}$ measurable and the future behavior of the arrivals which is a renewal process. Hence, we have

Theorem 7. *Longest-Idle-Server-First policy is totally blind and satisfies (4.2) and hence the limiting fairness process is as given in Theorem 6.*

Proof. For any fixed $T > 0$, $\sup_{0 \leq t \leq T} \mathbb{E}[u_{A^n(s)+1}^n|\mathcal{F}_t]$ is bounded and $\left\{ |\hat{I}^n|_{*,T} \right\}_{n \in \mathbb{N}}$ is tight, which implies the (4.2) holds. Also, for any i and n , (4.4) implies $\mathbb{E}[\phi_i^n|\mathcal{F}_{\theta_i^n}] - \mathbb{E}[\phi_i^n|\mathcal{F}_{\theta_i^n-}] = 0$ and hence LISF policy is totally blind. \square

5 Conclusion

In this work, we analyze many-server queueing systems where the service rates of servers in the system are i.i.d. random variables. These systems have been studied in an earlier paper by Atar [2] for two routing policies, longest-idle-server-first and fastest-server-first policies, in an ad hoc manner. Our main contribution is to develop a general framework to analyze many-server queueing systems with service rate uncertainty using measure-valued stochastic processes. We introduce the fairness process which assumes values in the space of probability measures denoting the proportion of cumulative idleness shared by servers having different service rates. Unfortunately, it is possible to show that the fairness processes cannot be analyzed in the standard Skorokhod topologies and we introduce a modified notion of convergence for these processes. We also include customer abandonments into our analysis, which we believe to be indispensable for systems with service rate uncertainty. We also show how the limits obtained in [2] can be obtained using martingale methods.

Even though, the many-server systems with uncertain service rates can be seen as a replacement for servers belonging to pools where servers are identical, we believe that this work also should be extended to include networks of queues with pools of servers, where within a server pool the service rates follow the same distribution, but may follow different distributions among pools. The introduced diffusion limits can also be used to analyze staffing of many-server systems with service rate uncertainty, and how the variability of the service rate affect the staffing and routing decisions. We also believe that the similar fairness processes can be defined to analyze service systems with heterogeneous and non-exponential service times and data-driven systems where the service rates are either time dependent or learned through data as time evolves.

Acknowledgements

We would like to thank Istvan Gyöngi, Sandy Davie and Gonçalo Dos Reis for the valuable discussions on this problem.

References

- [1] M. Armony. Dynamic routing in large-scale service systems with heterogeneous servers. *Queueing Systems*, 51:287–329, 2005.
- [2] R. Atar. Central limit theorem for a many-server queue with random service rates. *The Annals of Applied Probability*, 18(4):1548–1568, 2008.
- [3] R. Atar, Y. Shaki, and A. Schwartz. A blind policy for equalizing cumulative idleness. *Queueing Systems*, 67:275–293, 2011.
- [4] R. Atar and A. Schwartz. Efficient routing in heavy traffic under partial sampling of service times. *Mathematics of Operations Research*, 33(4):899–909, 2008.
- [5] A. Bassamboo, R. S. Randhawa, and A. Zeevi. Capacity sizing under parameter uncertainty: Safety staffing principles revisited. 56:1668–1686, 2010.
- [6] P. Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1999.

- [7] L. Decreusefond and P. Moyal. Fluid limit of a heavily loaded EDF queue with impatient customers. *Markov Processes and Related Fields*, 14:131–158, 2008.
- [8] N. Gans, N. Liu, A. Mandelbaum, H. Shen, and H. Ye. Service times in call centers: Agent heterogeneity and learning with some operational consequences. *Borrowing Strength: Theory Powering Applications* A Festschrift for Lawrence D. Brown, pages 99–123. Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2010.
- [9] O. Garnett, A. Mandelbaum, and M. Reiman. Designing a call center with impatient customers. *Manufacturing & Service Operations Management*, 4(3):208–227, 2002.
- [10] H. Gromoll, A. Puha, and R. Williams. The fluid limit of a heavily loaded processor sharing queue. *Annals of Applied Probability*, 12:797–859, 2002.
- [11] H. Gromoll, P. Robert, and B. Zwart. Fluid limits for processor-sharing queues with impatience. *Mathematics of Operations Research*, 33:375–402, 2008.
- [12] I. Gurvich and W. Whitt. Queue-and-idleness-ratio controls in many-server service systems. *Mathematics of Operations Research*, 34:363–396, 2009.
- [13] S. Halfin and W. Whitt. Heavy traffic limits for queues with many exponential servers. *Operations Research*, 29:567–588, 1981.
- [14] J. M. Harrison and A. Zeevi. A method for staffing large call centers using stochastic fluid models. *Manufacturing & Service Operations Management*, 7:20–36, 2005.
- [15] R. Ibrahim. Managing queueing systems where capacity is random and customers are impatient. forthcoming in *Production and Operations Management*.
- [16] A. Jakubowski. On the Skorokhod topology. *Annales de l’institut Henri Poincaré (B) Probability and Statistics*, 22(3):263–285, 1986.
- [17] W. Kang and K. Ramanan. Fluid limits of many server queues with reneging. *Annals of Applied Probability*, 20(6):2204–2260, 2010.
- [18] H. Kaspi and K. Ramanan. Law of large numbers limits for many-server queues. *Annals of Applied Probability*, 21(1):33–114, 2011.
- [19] H. Kaspi and K. Ramanan. SPDE limits of many-server queues. *Annals of Applied Probability*, 23(1):145–229, 2013.
- [20] I. Mitoma. Tightness of Probabilities On $C([0, 1], Y')$ and $D([0, 1], Y')$. *The Annals of Probability*, 11(4):989–999, 1983.
- [21] J. Munkres. *Topology*. Prentice Hall, Inc, New Jersey, second edition, 2000.
- [22] G. Pang, R. Talreja, and W. Whitt. Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probability Surveys*, 4:193–267, 2007.
- [23] P. Protter. *Stochastic Integration and Differential Equations*. Stochastic Modelling and Applied Probability. Springer, Berlin Heidelberg, second edition, 2005.

- [24] J. Reed and R. Talreja. Distribution-valued heavy-traffic limits for the $g/gi/\infty$ queue. *Annals of Applied Probability*, 25(3):1420–1474, 2015.
- [25] A. Skorokhod. Limit theorems for stochastic processes. *Theory of Probability and Its Applications*, 1(3):261–290, 1956.
- [26] A. Ward and M. Armony. Blind fair routing in large-scale service systems with heterogeneous customers and servers. *Operations Research*, 1(61):228–243, 2013.
- [27] W. Whitt. Weak convergence of probability measures on the function space $C[0, \infty)$. *The Annals of Mathematical Statistics*, 41(3):939–944, 1970.
- [28] W. Whitt. Some useful functions for functional limit theorems. *Mathematics of Operations Research*, 5(1):67–85, 1980.
- [29] J. Zan, J. Hasenbein, and D. Morton. Staffing call centers under arrival-rate uncertainty with Bayesian updates. *Operations Research Letters*, 46:379–384, 2018.