

Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Universality of Power-of-d Load Balancing in Many-Server Systems

Debankur Mukherjee, Sem C. Borst, Johan S. H. van Leeuwen, Philip A. Whiting

To cite this article:

Debankur Mukherjee, Sem C. Borst, Johan S. H. van Leeuwen, Philip A. Whiting (2018) Universality of Power-of-d Load Balancing in Many-Server Systems. Stochastic Systems 8(4):265-292. <https://doi.org/10.1287/stsy.2018.0016>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Universality of Power-of- d Load Balancing in Many-Server Systems

 Debankur Mukherjee,^a Sem C. Borst,^{a,b} Johan S. H. van Leeuwen,^a Philip A. Whiting^c
^aEindhoven University of Technology, 5600 MB Eindhoven, Netherlands; ^bNokia Bell Labs, Murray Hill, New Jersey 07974;

^cMacquarie University, North Ryde, New South Wales 2109, Australia

Contact: debankur_mukherjee@brown.edu,  <http://orcid.org/0000-0003-1678-4893> (DM); s.c.borst@tue.nl (SCB); j.s.h.v.leeuwen@tue.nl (JSHvL); philip.whiting@mq.edu.au (PAW)

Received: September 5, 2017


Revised: February 25, 2018

Accepted: May 26, 2018

Published Online in Articles in Advance: December 27, 2018

<https://doi.org/10.1287/stsy.2018.0016>
Copyright: © 2018 The Author(s)

Abstract. We consider a system of N parallel single-server queues with unit exponential service rates and a single dispatcher where tasks arrive as a Poisson process of rate $\lambda(N)$. When a task arrives, the dispatcher assigns it to a server with the shortest queue among $d(N)$ randomly selected servers ($1 \leq d(N) \leq N$). This load balancing strategy is referred to as a JSQ($d(N)$) scheme, noting that it subsumes the celebrated Join-the-Shortest Queue (JSQ) policy as a crucial special case for $d(N) = N$. We construct a stochastic coupling to bound the difference in the queue length processes between the JSQ policy and a JSQ($d(N)$) scheme with an arbitrary value of $d(N)$. We use the coupling to derive the fluid limit in the regime where $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$ with $d(N) \rightarrow \infty$, along with the associated fixed point. The fluid limit turns out not to depend on the exact growth rate of $d(N)$ and in particular coincides with that for the JSQ policy. We further leverage the coupling to establish that the diffusion limit in the critical regime where $(N - \lambda(N))/\sqrt{N} \rightarrow \beta > 0$ as $N \rightarrow \infty$ with $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$ corresponds to that for the JSQ policy. These results indicate that the optimality of the JSQ policy can be preserved at the fluid level and diffusion level while reducing the overhead by nearly a factor $O(N)$ and $O(\sqrt{N}/\log(N))$, respectively.

 **Open Access Statement:** This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Stochastic Systems*. Copyright © 2018 The Author(s). <https://doi.org/10.1287/stsy.2018.0016>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: This research was financially supported by an ERC Starting Grant and by the Netherlands Organization for Scientific Research (NWO) through TOP-GO [Grant 613.001.012] and Gravitation Networks [Grant 024.002.003]. P. A. Whiting was supported in part by the Australian Research Council [Grant DP-1592400] and in part by a Macquarie University Vice-Chancellor Innovation Fellowship.

Keywords: load balancing • power-of- d scheme • join the shortest queue • stochastic coupling • functional limit theorems • fluid limit • diffusion limit • many-server asymptotics

1. Introduction

In this paper, we establish a universality property for a broad class of randomized load balancing schemes in many-server systems. Whereas the specific features of load balancing policies may considerably differ, the principal purpose is to distribute service requests or tasks among servers or distributed resources in parallel-processing systems. Well-designed load balancing schemes provide an effective mechanism for improving relevant performance metrics experienced by users while achieving high resource utilization levels. The analysis and design of load balancing schemes has attracted strong renewed interest in the last several years, mainly motivated by significant challenges involved in assigning tasks (e.g., file transfers, compute jobs, and database look-ups) to servers in large-scale data centers.

In the present paper, we focus on a basic scenario of sending tasks from a single dispatcher to N parallel queues with identical servers, exponentially distributed service requirements, and a service discipline at each individual server that is oblivious to the actual service requirements (e.g., first-come first-served). In this canonical case, the so-called Join-the-Shortest-Queue (JSQ) policy has several strong optimality properties and, in particular, minimizes the overall mean delay among the class of nonanticipating load balancing policies that do not have any advance knowledge of the service requirements (Winston 1977, Weber 1978, Ephremides et al. 1980). (Relaxing any of the three above-mentioned assumptions tends to break the optimality properties of the JSQ policy and renders the delay-minimizing policy quite complex or even counterintuitive (see, for instance, Whitt 1986, Jonckheere 2006, and Gupta et al. 2007).

In order to implement the JSQ policy, a dispatcher requires instantaneous knowledge of the queue lengths at all the servers, which may give rise to a substantial communication burden and not be scalable in scenarios with large numbers of servers. The latter issue has motivated consideration of so-called JSQ(d) strategies, where the dispatcher

assigns an incoming task to a server with the shortest queue among d servers selected uniformly at random. Mean-field limit theorems by Mitzenmacher (2001) and Vvedenskaya et al. (1996) indicate that even a value as small as $d = 2$ yields significant performance improvements in a many-server regime with $N \rightarrow \infty$, in the sense that the tail of the queue length distribution at each individual server falls off much more rapidly compared with a strictly random assignment policy ($d = 1$). This is commonly referred to as the “power-of-two” effect. While these results were originally proved for exponential service requirement distributions, they have been extended to general service requirement distributions by Bramson et al. (2012). Analyses of several variants of this model can be found in the works of Graham (2005), Luczak and Norris (2005), Luczak and McDiarmid (2006), Brightwell et al. (2018), and Eschenfeldt and Gamarnik (2016). We refer to Van der Boor et al. (2018) for a recent survey on scalable load balancing algorithms.

The diversity parameter d thus induces a fundamental trade-off between the amount of communication overhead and the performance in terms of queue lengths and delays. Specifically, a strictly random assignment policy can be implemented with zero overhead, but for any positive load per server, the probability of nonzero wait and the mean waiting time do *not* fall to zero as $N \rightarrow \infty$. In contrast, a nominal implementation of the JSQ policy (without maintaining state information at the dispatcher) involves $O(N)$ overhead per task, but it can be shown that the probability of nonzero wait and the mean waiting time vanish as $N \rightarrow \infty$ for any fixed subcritical load per server. Although JSQ(d) strategies with a fixed parameter $d \geq 2$ yield significant performance improvements over purely random task assignment while reducing the communication overhead by a factor $O(N)$ compared with the JSQ policy, the probability of nonzero wait and mean waiting time do *not* vanish in the limit. In that sense a fixed value of d is not sufficient to achieve asymptotically optimal performance. This is also reflected by recent results of Gamarnik et al. (2016) indicating that in the absence of any memory at the dispatcher the communication overhead per task must grow with N in order to allow a zero mean waiting time in the limit.

In order to gain further insight in the trade-off between performance and communication overhead as governed by the diversity parameter d , we also consider a regime where the number of servers N grows large but allow the value of d to depend on N and write $d(N)$ to explicitly reflect that. For convenience, we assume a Poisson arrival process of rate $\lambda(N)$ and unit-mean exponential service requirements.

We construct a stochastic coupling to bound the difference in the queue length processes between the ordinary JSQ policy and a scheme with an arbitrary value of $d(N)$. We exploit the coupling to obtain the fluid limit in the subcritical regime where $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$ with $d(N) \rightarrow \infty$, along with the associated fixed point. As it turns out, the fluid limit does not depend on the exact growth rate of $d(N)$ and, in particular, coincides with that for the JSQ policy. This implies that the overhead of the JSQ policy can be reduced by “almost” a factor $O(N)$ while maintaining fluid-level optimality. In case of batch arrivals fluid-level optimality can even be achieved with $O(1)$ communication overhead per task.

We further consider the Halfin-Whitt heavy-traffic regime where $(N - \lambda(N))/\sqrt{N} \rightarrow \beta > 0$ as $N \rightarrow \infty$. Recent work of Eschenfeldt and Gamarnik (2018) showed that the diffusion-scaled system occupancy state for the ordinary JSQ policy in this regime weakly converges to a two-dimensional reflected Ornstein-Uhlenbeck process. We leverage the above-mentioned coupling to prove that the diffusion limit in case $d(N)/(\sqrt{N} \log N) \rightarrow \infty$ as $N \rightarrow \infty$ corresponds to that for the JSQ policy. This indicates that the overhead of the JSQ policy can “almost” be reduced to $O(\sqrt{N} \log N)$ while retaining diffusion-level optimality. The above condition is, in fact, close to necessary, in the sense that the diffusion-level behavior of the scheme is suboptimal if $d(N)/(\sqrt{N} \log N) \rightarrow 0$ as $N \rightarrow \infty$.

The above results mirror the fluid- and diffusion-level optimality properties reported in the companion paper Mukherjee et al. (2016a) for power-of- $d(N)$ strategies in a scenario with N server pools, where each server pool is a collection of servers, each working at unit rate. The coupling developed by Mukherjee et al. (2016a) has greater hold on the task completions and provides absolute bounds on the difference of each component of the occupancy states. More specifically, the task completions depend on the total number of active tasks in the entire system, whereas in the single-server scenario, it depends only on the number of nonidle servers. As a result, obtaining a stochastic coupling bound in this paper becomes analytically more challenging, and in contrast with the infinite-server scenario, involves the cumulative loss terms and tail sums of the occupancy states of the ordinary JSQ policy. This imposes the additional challenge of proving the ℓ_1 convergence of the occupancy state process of the ordinary JSQ policy as will be described in greater detail later. To the best of our knowledge, this is the first time the transient fluid limit of the ordinary JSQ policy is rigorously established.

The idea of using coupling to prove scaling limits of large-scale parallel-server systems was introduced by Mukherjee et al. (2016b). The coupling method there was much weaker and was useful only for systems starting from specific initial occupancy states and for the particular scaling regime considered in that paper. In contrast, in the current paper we need to develop a much stronger and wider coupling framework involving an intermediate class of schemes as described in Section 2.4 to establish the universality results. In addition, we

consider arbitrary starting states and different scaling regimes. Remark 1 further discusses the novelty and importance of the current stochastic comparison framework.

The remainder of the paper is organized as follows. In Section 2, we present a detailed model description and state the main results; in Section 3, we construct a coupling and establish the stochastic ordering relations. Sections 4 and 5 contain the proofs of the fluid and diffusion limit results, respectively. Finally, in Section 6, we make some concluding remarks and briefly comment on future research directions.

2. Main Results

2.1. Model Description and Notation

Consider a system with N parallel single-server queues with identical servers and a single dispatcher. Tasks with unit-mean exponential service requirements arrive at the dispatcher as a Poisson process of rate $\lambda(N)$ and are instantaneously forwarded to one of the servers. Specifically, when a task arrives, the dispatcher assigns it to a server with the shortest queue among $d(N)$ randomly selected servers ($1 \leq d(N) \leq N$). This load balancing strategy will be referred to as the JSQ($d(N)$) scheme, noting that it subsumes the ordinary JSQ policy as a crucial special case for $d(N) = N$. The buffer capacity at each of the servers is b (possibly infinite), and when a task is assigned to a server with b pending tasks, it is permanently discarded.

For any $d(N)$ ($1 \leq d(N) \leq N$), let

$$\mathbf{Q}^{d(N)}(t) := \left(Q_1^{d(N)}(t), Q_2^{d(N)}(t), \dots, Q_b^{d(N)}(t) \right)$$

denote the system occupancy state, where $Q_i^{d(N)}(t)$ is the number of servers under the JSQ($d(N)$) scheme with a queue length of i or larger, at time t , including the possible task in service, $i = 1, \dots, b$. Figure 1 provides a schematic diagram of the Q_i values. Throughout we assume that at each arrival epoch the servers are ordered in nondecreasing order of their queue lengths (ties can be broken arbitrarily), and whenever we refer to some ordered server, it should be understood to be with respect to this prior ordering.

We occasionally omit the superscript $d(N)$ and replace it by N to refer to the N th system, when the value of $d(N)$ is clear from the context. When a task is discarded, in case of a finite buffer size, we call it an *overflow* event, and we denote by $L^{d(N)}(t)$ the total number of overflow events under the JSQ($d(N)$) scheme up to time t .


A sequence of random variables $\{X_N\}_{N \geq 1}$, for some function $f : \mathbb{R} \rightarrow \mathbb{R}_+$, is said to be $O_P(f(N))$ if the sequence of scaled random variables $\{X_N/f(N)\}_{N \geq 1}$ is tight or said to be $o_P(f(N))$ if $\{X_N/f(N)\}_{N \geq 1}$ converges to zero in probability. Boldfaced letters are used to denote vectors. We denote by ℓ_1 the space of all summable sequences. For any set K , the closure is denoted by \bar{K} . We denote by $D_E[0, \infty)$ the set of all *càdlàg* (right continuous left limit exists) functions from $[0, \infty)$ to a complete separable metric space E and by " \xrightarrow{d} " convergence in distribution for real-valued random variables and with respect to the Skorohod- J_1 topology for càdlàg processes.

2.2. Fluid-Limit Results

In the fluid-level analysis, we consider the subcritical regime where $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$. In order to state the results, we first introduce some useful notation. Denote the fluid-scaled system occupancy state by $\mathbf{q}^{d(N)}(t) := \mathbf{Q}^{d(N)}(t)/N$, i.e., $q_i^{d(N)}(t) = Q_i^{d(N)}(t)/N$, and define

$$\mathcal{S} = \left\{ \mathbf{q} \in [0, 1]^b : q_i \leq q_{i-1} \text{ for all } i = 2, \dots, b, \text{ and } \sum_{i=1}^b q_i < \infty \right\}$$

Figure 1. Occupancy State of the System



Note. When the servers are arranged in nondecreasing order of their queue lengths, Q_i represents the width of the i th row.

as the set of all possible fluid-scaled occupancy states equipped with the ℓ_1 topology. For any $\mathbf{q} \in \mathcal{S}$, denote $m(\mathbf{q}) = \min\{i : q_{i+1} < 1\}$, with the convention that $q_{b+1} = 0$ if $b < \infty$. Note that $m(\mathbf{q}) < \infty$ because $\mathbf{q} \in \ell_1$. If $m(\mathbf{q}) = 0$, then define $p_0(\mathbf{q}) = 1$ and $p_i(\mathbf{q}) = 0$ for all $i \geq 1$. If $m(\mathbf{q}) > 0$, distinguish two cases, depending on whether the normalized arrival rate λ is larger than $1 - q_{m(\mathbf{q})+1}$ or not. If $\lambda < 1 - q_{m(\mathbf{q})+1}$, then define $p_{m(\mathbf{q})-1}(\mathbf{q}) = 1$ and $p_i(\mathbf{q}) = 0$ for all $i \neq m(\mathbf{q}) - 1$. On the other hand, if $\lambda > 1 - q_{m(\mathbf{q})+1}$, then $p_{m(\mathbf{q})-1}(\mathbf{q}) = (1 - q_{m(\mathbf{q})+1})/\lambda$, $p_{m(\mathbf{q})}(\mathbf{q}) = 1 - p_{m(\mathbf{q})-1}(\mathbf{q})$, and $p_i(\mathbf{q}) = 0$ for all $i \neq m(\mathbf{q}) - 1, m(\mathbf{q})$. Note that the assumption $\lambda < 1$ ensures that the latter case cannot occur when $m(\mathbf{q}) = b < \infty$.

Theorem 1 (Universality of Fluid Limit for JSQ($d(N)$) Scheme). *Assume $\mathbf{q}^{d(N)}(0) \rightarrow \mathbf{q}^\infty$ in \mathcal{S} and $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$. For the JSQ($d(N)$) scheme with $d(N) \rightarrow \infty$, any subsequence of the sequence of processes $\{\mathbf{q}^{d(N)}(t)\}_{t \geq 0}$ has a further subsequence that converges weakly with respect to the Skorohod- J_1 topology, to the limit $\{\mathbf{q}(t)\}_{t \geq 0}$ satisfying the following system of integral equations*

$$q_i(t) = q_i^\infty + \lambda \int_0^t p_{i-1}(\mathbf{q}(s)) ds - \int_0^t (q_i(s) - q_{i+1}(s)) ds, \quad i = 1, \dots, b, \tag{1}$$

where the coefficients $p_i(\cdot)$ are as defined earlier.

The above theorem shows that the fluid-level dynamics do not depend on the specific growth rate of $d(N)$ as long as $d(N) \rightarrow \infty$ as $N \rightarrow \infty$. In particular, the JSQ($d(N)$) scheme with $d(N) \rightarrow \infty$ exhibits the same behavior as the ordinary JSQ policy in the limit, and thus achieves fluid-level optimality.

The coefficient $p_i(\mathbf{q})$ represents the instantaneous fraction of incoming tasks assigned to servers with a queue length of exactly i in the fluid-level state $\mathbf{q} \in \mathcal{S}$. Assuming $m(\mathbf{q}) < b$, a strictly positive fraction $1 - q_{m(\mathbf{q})+1}$ of the servers have a queue length of exactly $m(\mathbf{q})$. Because $d(N) \rightarrow \infty$, the fraction of incoming tasks that get assigned to servers with a queue length of $m(\mathbf{q}) + 1$ or larger is zero: $p_i(\mathbf{q}) = 0$ for all $i = m(\mathbf{q}) + 1, \dots, b - 1$. Also, tasks at servers with a queue length of exactly i are completed at (normalized) rate $q_i - q_{i+1}$, which is zero for all $i = 0, \dots, m(\mathbf{q}) - 1$, and hence the fraction of incoming tasks that get assigned to servers with a queue length of $m(\mathbf{q}) - 2$ or less is zero as well: $p_i(\mathbf{q}) = 0$ for all $i = 0, \dots, m(\mathbf{q}) - 2$. This only leaves the fractions $p_{m(\mathbf{q})-1}(\mathbf{q})$ and $p_{m(\mathbf{q})}(\mathbf{q})$ to be determined. Now, observe that the fraction of servers with a queue length of exactly $m(\mathbf{q}) - 1$ is zero. If $m(\mathbf{q}) = 0$, then clearly the incoming tasks will join the empty queue, and thus, $p_{m(\mathbf{q})} = 1$ and $p_i(\mathbf{q}) = 0$ for all $i \neq m(\mathbf{q})$. Furthermore, if $m(\mathbf{q}) \geq 1$, because tasks at servers with a queue length of exactly $m(\mathbf{q})$ are completed at (normalized) rate $1 - q_{m(\mathbf{q})+1} > 0$, incoming tasks can be assigned to servers with a queue length of exactly $m(\mathbf{q}) - 1$ at that rate. We thus need to distinguish between two cases, depending on whether the normalized arrival rate λ is larger than $1 - q_{m(\mathbf{q})+1}$ or not. If $\lambda < 1 - q_{m(\mathbf{q})+1}$, then all the incoming tasks can be assigned to a server with a queue length of exactly $m(\mathbf{q}) - 1$, so that $p_{m(\mathbf{q})-1}(\mathbf{q}) = 1$ and $p_{m(\mathbf{q})}(\mathbf{q}) = 0$. On the other hand, if $\lambda > 1 - q_{m(\mathbf{q})+1}$, then not all incoming tasks can be assigned to servers with a queue length of exactly $m(\mathbf{q}) - 1$ active tasks, and a positive fraction will be assigned to servers with a queue length of exactly $m(\mathbf{q})$: $p_{m(\mathbf{q})-1}(\mathbf{q}) = (1 - q_{m(\mathbf{q})+1})/\lambda$ and $p_{m(\mathbf{q})}(\mathbf{q}) = 1 - p_{m(\mathbf{q})-1}(\mathbf{q})$.

It is easily verified that the unique fixed point $\mathbf{q}^* = (q_1^*, q_2^*, \dots, q_b^*)$ of the system of differential equations in Equation (1) is given by

$$q_i^* = \begin{cases} \lambda, & i = 1, \\ 0, & i = 2, \dots, b. \end{cases} \tag{2}$$

Note that the fixed point in Equation (2) is consistent with the results of Mitzenmacher (2001), Vvedenskaya et al. (1996), and Ying et al. (2015) for fixed d , where taking $d \rightarrow \infty$ yields the same fixed point. However, the results of Mitzenmacher (2001), Vvedenskaya et al. (1996), and Ying et al. (2015) for fixed d cannot be directly used to handle joint scalings and do not yield the universality of the entire fluid-scaled sample path for arbitrary initial states as established in Theorem 1.

The fixed point in Equation (2) in conjunction with the interchange of limits result in Proposition 1 below indicates that in stationarity the fraction of servers with a queue length of two or larger is negligible. Let

$$\pi^{d(N)}(\cdot) = \lim_{t \rightarrow \infty} \mathbb{P}(\mathbf{q}^{d(N)}(t) = \cdot)$$

be the stationary measure of the occupancy states of the N th system.

Proposition 1 (Interchange of Limits). *For the JSQ($d(N)$) scheme let $\pi^{d(N)}$ be the stationary measure of the occupancy states of the N th system. Then, $\pi^{d(N)} \xrightarrow{d} \pi^*$ as $N \rightarrow \infty$ with $d(N) \rightarrow \infty$, where $\pi^* = \delta_{\mathbf{q}^*}$ with δ_x being the Dirac measure concentrated upon x , and \mathbf{q}^* as in Equation (2).*

The above proposition relies on tightness of $\{\pi^{d(N)}\}_{N \geq 1}$ and the global stability of the fixed point and is proved in Subsection 4.3.

We now consider an extension of the model in which tasks arrive in batches. We assume that the batches arrive as a Poisson process with rate $\lambda(N)/\ell(N)$ and have fixed size $\ell(N) > 0$, so that the effective total task arrival rate remains $\lambda(N)$. We will show that even for arbitrarily slowly growing batch size, fluid-level optimality can be achieved with $O(1)$ communication overhead per task. For that, we define the JSQ($d(N)$) scheme adapted for batch arrivals. When a batch of size $\ell(N)$ arrives, the dispatcher samples $d(N) \geq \ell(N)$ servers without replacement and assigns the $\ell(N)$ tasks to the $\ell(N)$ servers with the smallest queue length among the sampled servers.

Theorem 2 (Batch Arrivals). *Consider the batch arrival scenario with growing batch size $\ell(N) \rightarrow \infty$ and $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$. For the JSQ($d(N)$) scheme with $d(N) \geq \ell(N)/(1 - \lambda - \varepsilon)$ for any fixed $\varepsilon > 0$, if $q_1^{d(N)}(0) \xrightarrow{\mathbb{P}} q_1^\infty \leq \lambda$ and $q_i^{d(N)}(0) \xrightarrow{\mathbb{P}} 0$ for all $i \geq 2$, then the sequence of processes $\{\mathbf{q}^{d(N)}(t)\}_{t \geq 0}$ converges weakly to the limit $\{\mathbf{q}(t)\}_{t \geq 0}$ described as follows:*

$$q_1(t) = \lambda + (q_1^\infty - \lambda)e^{-t}, \quad q_i(t) \equiv 0 \quad \text{for all } i = 2, \dots, b. \quad (3)$$

The fluid limit in Equation (3) agrees with the fluid limit of the JSQ($d(N)$) scheme if the initial state is taken as in Theorem 2. Further observe that the fixed point also coincides with that of the JSQ policy, as given by Equation (2). Also, for a fixed $\varepsilon > 0$, the communication overhead per task is on average given by $(1 - \lambda - \varepsilon)^{-1}$, which is $O(1)$. Thus, Theorem 2 ensures that in case of batch arrivals with growing batch size, fluid-level optimality can be achieved with $O(1)$ communication overhead per task. The result for the fluid-level optimality in stationarity can also be obtained indirectly by exploiting the fluid-limit result of Ying et al. (2015). Specifically, it can be deduced from the result of Ying et al. (2015) that for batch arrivals with growing batch size, the JSQ($d(N)$) scheme with suitably growing $d(N)$ yields the same fixed point of the fluid limit as described in Equation (2).

2.3. Diffusion-Limit Results

In the diffusion-limit analysis, we consider the Halfin-Whitt regime where

$$\frac{N - \lambda(N)}{\sqrt{N}} \rightarrow \beta \quad \text{as } N \rightarrow \infty$$

for some positive coefficient $\beta > 0$. In order to state the results, we first introduce some useful notation. Let $\bar{\mathbf{Q}}^{d(N)}(t) = (\bar{Q}_1^{d(N)}(t), \bar{Q}_2^{d(N)}(t), \dots, \bar{Q}_b^{d(N)}(t))$ be a properly centered and scaled version of the system occupancy state $\mathbf{Q}^{d(N)}(t)$, with

$$\bar{Q}_1^{d(N)}(t) = -\frac{N - Q_1^{d(N)}(t)}{\sqrt{N}}, \quad \bar{Q}_i^{d(N)}(t) = \frac{Q_i^{d(N)}(t)}{\sqrt{N}}, \quad i = 2, \dots, b.$$

The reason why $Q_1^{d(N)}(t)$ is centered around N while $Q_i^{d(N)}(t)$, $i = 2, \dots, b$ are not is because the fraction of servers with a queue length of exactly one tends to one, whereas the fraction of servers with a queue length of two or more tends to zero as $N \rightarrow \infty$.

Theorem 3 (Universality of Diffusion Limit for JSQ($d(N)$) Scheme). *Assume $\bar{Q}_i^{d(N)}(0) \xrightarrow{d} \bar{Q}_i(0)$ in \mathbb{R} as $N \rightarrow \infty$, buffer capacity $b \geq 2$ (possibly infinite), and there exists some $k \geq 2$ such that $\bar{Q}_{k+1}^N(0) = 0$ for all sufficiently large N . For $d(N)/(\sqrt{N} \log N) \rightarrow \infty$, the sequence of processes $\{\bar{\mathbf{Q}}^{d(N)}(t)\}_{t \geq 0}$ converges weakly to the limit $\{\bar{\mathbf{Q}}(t)\}_{t \geq 0}$ in $D_{\mathcal{F}}[0, \infty)$, where $\bar{Q}_i(t) \equiv 0$ for $i \geq k + 1$ and $(\bar{Q}_1(t), \bar{Q}_2(t), \dots, \bar{Q}_k(t))$ are the unique solutions in $D_{\mathbb{R}^k}[0, \infty)$ of the stochastic integral equations*

$$\begin{aligned} \bar{Q}_1(t) &= \bar{Q}_1(0) + \sqrt{2}W(t) - \beta t + \int_0^t (-\bar{Q}_1(s) + \bar{Q}_2(s))ds - U_1(t), \\ \bar{Q}_2(t) &= \bar{Q}_2(0) + U_1(t) - \int_0^t (\bar{Q}_2(s) - \bar{Q}_3(s))ds, \\ \bar{Q}_i(t) &= \bar{Q}_i(0) - \int_0^t (\bar{Q}_i(s) - \bar{Q}_{i+1}(s))ds, \quad i = 3, \dots, k \end{aligned} \quad (4)$$

for $t \geq 0$, where W is the standard Brownian motion and U_1 is the unique nondecreasing nonnegative process in $D_{\mathbb{R}}[0, \infty)$ satisfying $\int_0^\infty \mathbf{1}_{[\bar{Q}_1(t) < 0]} dU_1(t) = 0$.

Although Equation (4) differs from the diffusion limit obtained for the fully pooled M/M/N model in the Halfin-Whitt regime Halfin and Whitt (1981) and Van Leeuwen and Knessl (2011, 2012), it shares similar favorable properties, see Banerjee and Mukherjee (2018) and Braverman (2018). Observe that $-\bar{Q}_1^{d(N)}$ is the scaled number of

vacant servers. Thus, Theorem 3 shows that over any finite time horizon, there will be $O_P(\sqrt{N})$ servers with queue length zero and $O_P(\sqrt{N})$ servers with a queue length larger than two, and hence all but $O_P(\sqrt{N})$ servers have a queue length of exactly one. This diffusion limit is proved by Eschenfeldt and Gamarnik (2018) for the ordinary JSQ policy. Our contribution is to construct a stochastic coupling and establish that, somewhat remarkably, the diffusion limit is the same for any JSQ($d(N)$) scheme, as long as $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$. In particular, the JSQ($d(N)$) scheme with $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$ exhibits the same behavior as the ordinary JSQ policy in the limit and thus achieves diffusion-level optimality. This growth condition for $d(N)$ is not only sufficient but also nearly necessary, as indicated by the next theorem.

Theorem 4 (Almost Necessary Condition). *Assume $\bar{Q}_i^{d(N)}(0) \xrightarrow{d} \bar{Q}_i(0)$ in \mathbb{R} as $N \rightarrow \infty$. If $d(N)/(\sqrt{N} \log N) \rightarrow 0$, then the diffusion limit of the JSQ($d(N)$) scheme differs from that of the JSQ policy.*

Theorem 4, in conjunction with Theorem 3, shows that $\sqrt{N} \log N$ is the minimal order of $d(N)$ for the JSQ($d(N)$) scheme to achieve diffusion-level optimality.

2.4. Proof Strategy

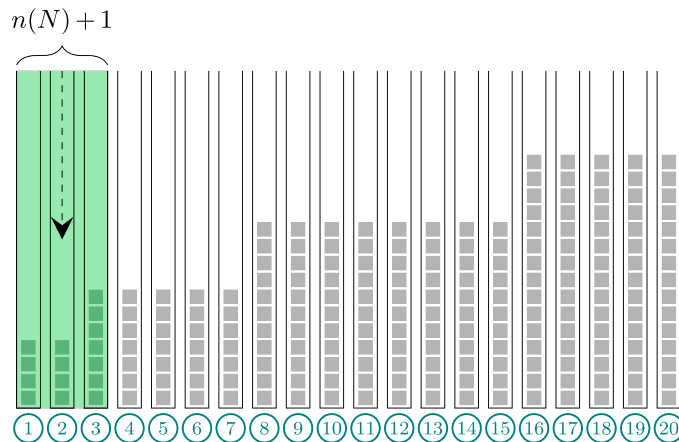
The idea behind the proofs of the asymptotic results for the JSQ($d(N)$) scheme in Theorems 1 and 3 is to (i) prove the fluid limit and exploit the existing diffusion limit result for the ordinary JSQ policy and then (ii) prove a universality result by establishing that the ordinary JSQ policy and the JSQ($d(N)$) scheme coincide under some suitable conditions on $d(N)$. For the ordinary JSQ policy the fluid limit in the subcritical regime is established in Subsection 4.1, and the diffusion limit in the Halfin-Whitt heavy-traffic regime by Eschenfeldt and Gamarnik (2018, theorem 2). A direct comparison between the JSQ($d(N)$) scheme and the ordinary JSQ policy is not straightforward, which is why we introduce the CJSQ($n(N)$) class of schemes as an intermediate scenario to establish the universality result.

Just like the JSQ($d(N)$) scheme, the schemes in the class CJSQ($n(N)$) may be thought of as “sloppy” versions of the JSQ policy, in the sense that tasks are not necessarily assigned to a server with the shortest queue length but to one of the $n(N) + 1$ lowest ordered servers, as graphically illustrated in Figure 2. In particular, for $n(N) = 0$, the class only includes the ordinary JSQ policy.

Note that the JSQ($d(N)$) scheme is guaranteed to identify the lowest ordered server but only among a randomly sampled subset of $d(N)$ servers. In contrast, a scheme in the CJSQ($n(N)$) class only guarantees that one of the $n(N) + 1$ lowest ordered servers is selected but across the entire pool of N servers. We will show that for sufficiently small $n(N)$, any scheme from the class CJSQ($n(N)$) is still “close” to the ordinary JSQ policy. We will further prove that for sufficiently large $d(N)$ relative to $n(N)$ we can construct a scheme called JSQ($n(N), d(N)$), belonging to the CJSQ($n(N)$) class, which differs ‘negligibly’ from the JSQ($d(N)$) scheme. Therefore, for a “suitable” choice of $d(N)$ the idea is to produce a “suitable” $n(N)$. This proof strategy is schematically represented in Figure 3.

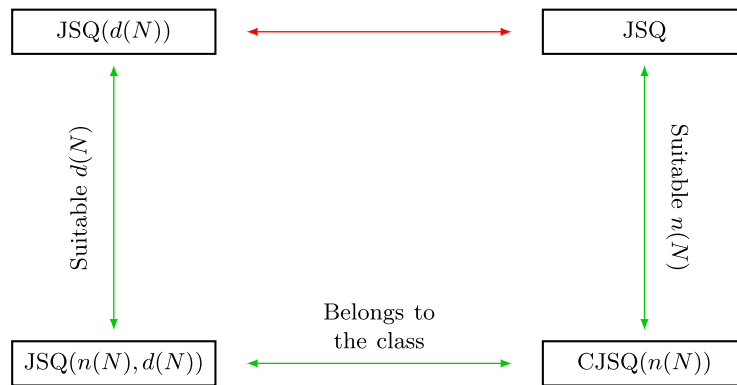
Remark 1. As mentioned in the introduction, a coupling method was used by Mukherjee et al. (2016b) to establish the diffusion limit of the Join-the-Idle Queue (JIQ) policy starting from specific initial occupancy states. Comparing the JIQ and JSQ policies in that scaling regime was much easier when viewed as follows: (i) If there is an idle server

Figure 2. CJSQ($n(N)$) Scheme



Note. High-level view of the CJSQ($n(N)$) class of schemes, where, as in Figure 1, the servers are arranged in the nondecreasing order of their queue lengths and the arrival must be assigned through the left tunnel.

Figure 3. Asymptotic Equivalence Relations



Note. The equivalence structure is depicted for various intermediate load balancing schemes to facilitate the comparison between the $JSQ(d(N))$ scheme and the ordinary JSQ policy.

in the system, both JIQ and JSQ perform similarly; (ii) also, when there is no idle server and only $O(\sqrt{N})$ servers with queue length two, JSQ assigns the arriving task to a server with queue length one. In that case, because JIQ assigns at random, the probability that the task will land on a server with queue length two and thus acts differently than JSQ is $O(1/\sqrt{N})$. Because on any finite time interval the number of times an arrival finds all servers busy is at most $O(\sqrt{N})$, all the arrivals except an $O(1)$ of them are assigned in exactly the same manner in both JIQ and JSQ, which then leads to the same scaling limit for both policies. Note that in the computation of the expected number of events when JIQ and JSQ act differently, both the specific initial state and the scaling regime were crucial. In the current paper the stochastic comparison framework is inherently different. Here the idea pivots on two key observations: (i) For any scheme, if each arrival is assigned to *approximately* the shortest queue, then the scheme can still retain its optimality on various scales, and (ii) for any two schemes, if on any finite time interval not *too many* arrivals are assigned to different ordered servers, then they can have the same scaling limits. Combining the above two ideas provides a much wider coupling framework involving an intermediate class of schemes that enables us to consider arbitrary starting states and different scaling regimes. In addition, the consideration of the arbitrary starting state will turn out to be crucial in order to extend the fluid-scale universality result to the steady state.

In the next section, we construct a stochastic coupling called S-coupling, which will be the key vehicle in establishing the universality result mentioned above.

Remark 2. Observe that, sampling *without* replacement polls more servers than *with* replacement; hence, the minimum number of active tasks among the selected servers is stochastically smaller in the case without replacement. As a result, for sufficient conditions as in Theorems 1 and 4, it is enough to consider sampling with replacement. Also, for notational convenience, in the proof of the almost necessary condition stated in Theorem 4 we will assume sampling with replacement, although the proof technique and the result are also valid if the servers are chosen without replacement.

3. Coupling and Stochastic Ordering

In this section, we construct a coupling between any scheme from the class $CJSQ(n(N))$ and the ordinary JSQ policy, which ensures that for sufficiently small $n(N)$, on any finite time interval, the two schemes differ negligibly. This plays an instrumental role in establishing the universality results in Theorems 1 and 3. All the statements in this section should be understood to apply to the N th system with N servers.

3.1. Stack Formation and Deterministic Ordering

In order to prove the stochastic comparisons among the various schemes, as in Mukherjee et al. (2016b), we describe the many-server system as an ensemble of stacks, in a way that two different ensembles can be ordered. In this formulation, at each step, items are added or removed according to some rule. From a high level, we then show that if two systems follow some specific rules, then at any step, the two ensembles maintain some kind of deterministic ordering. This deterministic ordering turns into an almost sure ordering in the next subsection, when we construct the S-coupling.

Each server along with its queue is thought of as a stack of items, and we always consider the stacks to be arranged in nondecreasing order of their heights. The ensemble of stacks then represents the empirical CDF of the queue length distribution, and the i th horizontal bar corresponds to Q_i^Π (for some task assignment scheme Π), as

depicted in Figure 1. If an arriving item happens to land on a stack that already contains b items, then the item is discarded, and is added to a special stack L^Π of discarded items, where it stays forever.

Any two ensembles A and B , each having N stacks and a maximum height b per stack, are said to follow Rule (n_A, n_B, k) at some step, if either an item is removed from the k th stack in both ensembles (if nonempty) or an item is added to the n_A th stack in ensemble A and to the n_B th stack in ensemble B .

Proposition 2. *For any two ensembles of stacks A and B , as described above, if at any step Rule (n_A, n_B, k) is followed for some value of n_A, n_B , and k , with $n_A \leq n_B$, then the following ordering is always preserved: for all $m \leq b$,*

$$\sum_{i=m}^b Q_i^A + L^A \leq \sum_{i=m}^b Q_i^B + L^B. \tag{5}$$

This proposition says that, while adding the items to the ordered stacks, if we ensure that in ensemble A the item is always placed to the left of that in ensemble B and if the items are removed from the same ordered stack in both ensembles, then the aggregate size of the $b - m + 1$ highest horizontal bars as depicted in Figure 1 plus the cumulative number of discarded items is no larger in A than in B throughout.

Proof of Proposition 2. We prove the ordering by forward induction on the time steps, i.e., we assume that at some step the ordering holds, and show that in the next step it will be preserved. In ensemble Π , where $\Pi = A, B$, after applying Rule (n_A, n_B, k) , the updated lengths of the horizontal bars are denoted by $\tilde{Q}_i^\Pi, i \geq 1$. Also, define $I_\Pi(c) := \max \{i \geq 0 : Q_i^\Pi \geq N - c + 1\}, c = 1, \dots, N$, with the convention that $Q_0^\Pi \equiv N$.

Now, if the rule prescribes removal of an item from the k th stack, then the updated ensemble will have the values

$$\tilde{Q}_i^\Pi = \begin{cases} Q_i^\Pi - 1, & \text{for } i = I_\Pi(k), \\ Q_j^\Pi, & \text{otherwise,} \end{cases} \tag{6}$$

if $I_\Pi(k) \geq 1$; otherwise, all the Q_i^Π -values remain unchanged. On the other hand, if the rule produces the addition of an item to stack n_Π , then the values will be updated as

$$\tilde{Q}_i^\Pi = \begin{cases} Q_i^\Pi + 1, & \text{for } i = I_\Pi(n_\Pi) + 1, \\ Q_j^\Pi, & \text{otherwise,} \end{cases} \tag{7}$$

if $I_\Pi(n_\Pi) < b_\Pi$; otherwise, all values remain unchanged.

Fix any $m \leq b$. Observe that in any event the Q_i -values change by at most one at any step, and hence it suffices to prove the preservation of the ordering in the case when Equation (5) holds with equality:

$$\sum_{i=m}^b Q_i^A + L^A = \sum_{i=m}^b Q_i^B + L^B. \tag{8}$$

We distinguish between two cases depending on whether an item is removed or added. First, suppose that the rule prescribes removal of an item from the k th stack from both ensembles. Observe from Equation (6) that the value of $\sum_{i=m}^b Q_i^\Pi + L^\Pi$ changes if and only if $I_\Pi(k) \geq m$. Also, because removal of an item can only decrease the sum, without loss of generality, we may assume that $I_B(k) \geq m$; otherwise, the right side of Equation (8) remains unchanged, and the ordering is trivially preserved. From our initial hypothesis,

$$\sum_{i=m+1}^b Q_i^A + L^A \leq \sum_{i=m+1}^b Q_i^B + L^B. \tag{9}$$

This implies

$$Q_m^A = \sum_{i=m}^b Q_i^A - \sum_{i=m+1}^b Q_i^A \geq \sum_{i=m}^b Q_i^B - \sum_{i=m+1}^b Q_i^B = Q_m^B. \tag{10}$$

Also,

$$\begin{aligned} I_B(k) \geq m &\iff Q_m^B \geq N - k + 1 \\ &\implies Q_m^A \geq N - k + 1 \iff I_A(k) \geq m. \end{aligned} \tag{11}$$

Therefore, the sum $\sum_{i=m}^b Q_i^A + L^A$ also decreases, and the ordering is preserved.

Now, suppose that the rule prescribes addition of an item to the respective stacks in both ensembles. From Equation (7) we get that after adding an item, the value of $\sum_{i=m}^b Q_i^\Pi + L^\Pi$ increases only if $I_\Pi(n_\Pi) \geq m - 1$. As in the previous case, we assume Equation (8), and because adding an item can only increase the concerned sums, we assume that $I_A(n_A) \geq m - 1$, because otherwise the left side of Equation (8) remains unchanged, and the ordering is trivially preserved. Now from our initial hypothesis we have

$$\sum_{i=m-1}^b Q_i^A + L^A \leq \sum_{i=m-1}^b Q_i^B + L^B. \quad (12)$$

Combining Equation (8) with Equation (12) gives

$$Q_{m-1}^A = \left(\sum_{i=m-1}^b Q_i^A + L^A \right) - \left(\sum_{i=m}^b Q_i^A + L^A \right) \leq \left(\sum_{i=m-1}^b Q_i^B + L^B \right) - \left(\sum_{i=m}^b Q_i^B + L^B \right) = Q_{m-1}^B. \quad (13)$$

Observe that

$$\begin{aligned} I_A(n_A) \geq m - 1 &\iff Q_{m-1}^A \geq N - n_A + 1 \implies Q_{m-1}^A \geq N - n_B + 1 \\ &\implies Q_{m-1}^B \geq N - n_B + 1 \iff I_B(n_B) \geq m - 1. \end{aligned} \quad (14)$$

Hence, the value of $\sum_{i=m}^b Q_i^B + L^B$ also increases, and the ordering is preserved. \square

3.2. Stochastic Ordering

We now use the deterministic ordering established in Proposition 2 in conjunction with the S-coupling construction to prove a stochastic comparison between the JSQ($d(N)$) scheme, a specific scheme from the class CJSQ($n(N)$) and the ordinary JSQ policy. As described earlier, the class CJSQ($n(N)$) contains all schemes that assign incoming tasks by some rule to any of the $n(N) + 1$ lowest ordered servers. Observe that when $n(N) = 0$, the class contains only the ordinary JSQ policy. Also, if $n^{(1)}(N) < n^{(2)}(N)$, then $\text{CJSQ}(n^{(1)}(N)) \subset \text{CJSQ}(n^{(2)}(N))$. Let MJSQ($n(N)$) be a particular scheme that always assigns incoming tasks to precisely the $(n(N) + 1)$ -th ordered server. Notice that this scheme is effectively the JSQ policy when the system always maintains $n(N)$ idle servers, or equivalently, uses only $N - n(N)$ servers, and $\text{MJSQ}(n(N)) \in \text{CJSQ}(n(N))$. For brevity, we suppress $n(N)$ in the notation for the remainder of this subsection.

We call any two systems *S-coupled*, if they have synchronized arrival clocks and departure clocks of the k th longest queue, for $1 \leq k \leq N$ (“S” in the name of the coupling stands for server). Consider three S-coupled systems following, respectively, the JSQ policy, any scheme from the class CJSQ and the MJSQ scheme. Recall that $Q_i^\Pi(t)$ is the number of servers with at least i tasks at time t and $L^\Pi(t)$ is the total number of lost tasks up to time t for the schemes $\Pi = \text{JSQ}, \text{CJSQ}, \text{and MJSQ}$. The following proposition provides a stochastic ordering for any scheme in the class CJSQ with respect to the ordinary JSQ policy and the MJSQ scheme.

Proposition 3. For any fixed $m \geq 1$,

- (i) $\left\{ \sum_{i=m}^b Q_i^{\text{JSQ}}(t) + L^{\text{JSQ}}(t) \right\}_{t \geq 0} \leq_{\text{st}} \left\{ \sum_{i=m}^b Q_i^{\text{CJSQ}}(t) + L^{\text{CJSQ}}(t) \right\}_{t \geq 0},$
- (ii) $\left\{ \sum_{i=m}^b Q_i^{\text{CJSQ}}(t) + L^{\text{CJSQ}}(t) \right\}_{t \geq 0} \leq_{\text{st}} \left\{ \sum_{i=m}^b Q_i^{\text{MJSQ}}(t) + L^{\text{MJSQ}}(t) \right\}_{t \geq 0},$

provided the inequalities hold at time $t = 0$.

The above proposition has the following immediate corollary, which will be used to prove bounds on the fluid and the diffusion scale.

Corollary 1. In the joint probability space constructed by the S-coupling of the three systems under, respectively, JSQ, MJSQ, and any scheme from the class CJSQ, the following ordering is preserved almost surely throughout the sample path: for any fixed $m \geq 1$,

- (i) $Q_m^{\text{CJSQ}}(t) \geq \sum_{i=m}^b Q_i^{\text{JSQ}}(t) - \sum_{i=m+1}^b Q_i^{\text{MJSQ}}(t) + L^{\text{JSQ}}(t) - L^{\text{MJSQ}}(t),$
- (ii) $Q_m^{\text{CJSQ}}(t) \leq \sum_{i=m}^b Q_i^{\text{MJSQ}}(t) - \sum_{i=m+1}^b Q_i^{\text{JSQ}}(t) + L^{\text{MJSQ}}(t) - L^{\text{JSQ}}(t),$

provided the inequalities hold at time $t = 0$.

Proof of Proposition 3. We first S-couple the concerned systems. Let us say that an incoming task is assigned to the n_Π th ordered server under scheme Π , $\Pi = \text{JSQ}, \text{CJSQ}, \text{MJSQ}$. Then, observe that, under the S-coupling, almost

surely, $n_{\text{JSQ}} \leq n_{\text{CJSQ}} \leq n_{\text{MJSQ}}$. Therefore, Proposition 2 ensures that in the probability space constructed through the S-coupling, the ordering is preserved almost surely throughout the sample path. \square

Remark 3. Note that $\sum_{i=1}^b \min\{Q_i, k\}$ represents the aggregate size of the rightmost k stacks, i.e., the k longest queues. Using this observation, the stochastic majorization property of the JSQ policy as stated by Towsley et al. (1992), Sparaggis et al. (1994), and Towsley (1995) can be shown following similar arguments as in the proof of Proposition 3. Conversely, the stochastic ordering between the JSQ policy and the MJSQ scheme presented in Proposition 3 can also be derived from the weak majorization arguments developed by Towsley et al. (1992), Sparaggis et al. (1994), and Towsley (1995). However, it is only through the stack arguments developed in the previous subsection that we could extend the results to compare any scheme from the class CJSQ with the scheme MJSQ as well as in Proposition 3(ii).

To analyze the JSQ($d(N)$) scheme, we need a further stochastic comparison argument. Consider two S-coupled systems following schemes Π_1 and Π_2 . Fix a specific arrival epoch, and let the arriving task join the n_{Π_i} th ordered server in the i th system following scheme Π_i , $i = 1, 2$ (ties can be broken arbitrarily in both systems). We say that at a specific arrival epoch the two systems *differ in decision* if $n_{\Pi_1} \neq n_{\Pi_2}$ and denote by $\Delta_{\Pi_1, \Pi_2}(t)$ the cumulative number of times the two systems differ in decision up to time t .

Proposition 4. For two S-coupled systems under schemes Π_1 and Π_2 the following inequality is preserved almost surely

$$\sum_{i=1}^b |Q_i^{\Pi_1}(t) - Q_i^{\Pi_2}(t)| \leq 2\Delta_{\Pi_1, \Pi_2}(t) \quad \forall t \geq 0, \tag{15}$$

provided the two systems start from the same occupancy state at $t = 0$, i.e., $Q_i^{\Pi_1}(0) = Q_i^{\Pi_2}(0)$ for all $i = 1, 2, \dots, b$.

Proof. We will again use forward induction on the event times of arrivals and departures. Let the inequality (15) hold at time epoch t_0 , and let t_1 be the next event time. We distinguish between two cases, depending on whether t_1 is an arrival epoch or a departure epoch.

If t_1 is an arrival epoch and the systems differ in decision, then observe that the left side of Equation (15) can only increase by two. In this case, the right side also increases by two, and the inequality is preserved. Therefore, it is enough to prove that the left side of Equation (15) remains unchanged if the two systems do not differ in decision. In that case, assume that both Π_1 and Π_2 assign the arriving task to the k th-ordered server. Recall from the proof of Proposition 2 the definition of I_{Π} for some scheme Π . If $I_{\Pi_1}(k) = I_{\Pi_2}(k)$, then the left side of Equation (15) clearly remains unchanged. Now, without loss of generality, assume $I_{\Pi_1}(k) < I_{\Pi_2}(k)$. Therefore,

$$Q_{I_{\Pi_1}(k)+1}^{\Pi_1}(t_0) < Q_{I_{\Pi_1}(k)+1}^{\Pi_2}(t_0).$$

After an arrival, the $(I_{\Pi_1}(k) + 1)$ -th term in the left side of Equation (15) decreases by one, and the $(I_{\Pi_2}(k) + 1)$ -th term may increase by at most one. Thus, the inequality is preserved.

If t_1 is a departure epoch, then due to the S-coupling, without loss of generality, assume that a potential departure occurs from the k th-ordered server. Also, note that a departure in either of the two systems can change at most one of the Q_i values. If at time epoch t_0 , $I_{\Pi_1}(k) = I_{\Pi_2}(k) = i$, then both $Q_i^{\Pi_1}$ and $Q_i^{\Pi_2}$ decrease by one, and hence the left side of Equation (15) does not change. Otherwise, without loss of generality assume $I_{\Pi_1}(k) < I_{\Pi_2}(k)$. Then, observe that

$$Q_{I_{\Pi_2}(k)}^{\Pi_1}(t_0) < Q_{I_{\Pi_2}(k)}^{\Pi_2}(t_0).$$

Furthermore, after the departure, $Q_{I_{\Pi_1}(k)}^{\Pi_1}$ may decrease at most by one. Therefore, $|Q_{I_{\Pi_1}(k)}^{\Pi_1} - Q_{I_{\Pi_1}(k)}^{\Pi_2}|$ may increase at most by one, and $Q_{I_{\Pi_2}(k)}^{\Pi_2}$ decreases by one; thus, $|Q_{I_{\Pi_2}(k)}^{\Pi_1} - Q_{I_{\Pi_2}(k)}^{\Pi_2}|$ decreases by one. Hence, in total, the left side of Equation (15) either remains the same or decreases by one. \square

3.3. Comparing the JSQ(d) and CJSQ(n) Schemes

We will now introduce the JSQ(n, d) scheme with $n, d \leq N$, which is an intermediate blend between the CJSQ(n) schemes and the JSQ(d) scheme. The JSQ(n, d) scheme will be seen in a moment to be a scheme in the CJSQ(n) class. It will also be seen to approximate the JSQ(d) scheme closely. We now specify the JSQ(d, n) scheme. At its first step, just as in the JSQ(d) scheme, it chooses the shortest of d random candidates but only sends this to that server's queue if it is one of the $n + 1$ shortest queues. Otherwise, it picks at the second step any of the $n + 1$ shortest

queues uniformly at random and then sends to that server’s queue. As was mentioned earlier, by construction, JSQ(d, n) is a scheme in CJSQ(n).

We now consider two S-coupled systems with a JSQ(d) and a JSQ(n, d) scheme. Assume that at some specific arrival epoch, the incoming task is dispatched to the k th ordered server in the system under the JSQ(d) scheme. If $k \in \{1, 2, \dots, n + 1\}$, then the system under JSQ(n, d) scheme also assigns the arriving task to the k th-ordered server. Otherwise, it dispatches the arriving task uniformly at random among the first $(n + 1)$ -ordered servers.

In the next proposition, we will bound the number of times these two systems differ in decision on any finite time interval. For any $T \geq 0$, let $A(T)$ and $\Delta(T)$ be the total number of arrivals to the system and the cumulative number of times that the JSQ(d) scheme and JSQ(n, d) scheme differ in decision up to time T .

Proposition 5. For any $T \geq 0$ and $M > 0$,

$$\mathbb{P}(\Delta(T) \geq M \mid A(T)) \leq \frac{A(T)}{M} \left(1 - \frac{n}{N}\right)^d. \quad (16)$$

Proof. Observe that at any arrival epoch, the systems under the JSQ(d) scheme and the JSQ(n, d) scheme will differ in decision only if none of the n lowest ordered servers gets selected by the JSQ(d) scheme. Now, at any arrival epoch, the probability that the JSQ(d) scheme does not select any of the n lowest ordered servers, is given by

$$p = \left(1 - \frac{n}{N}\right)^d.$$

Because at each arrival epoch, d servers are selected independently, given $A(T)$,

$$\Delta(T) \sim \text{Bin}(A(T), p).$$

Therefore, for $T \geq 0$, Markov’s inequality yields, for any fixed $M > 0$,

$$\mathbb{P}(\Delta(T) \geq M \mid A(T)) \leq \frac{\mathbb{E}(\Delta(T) \mid A(T))}{M} = \frac{A(T)}{M} \left(1 - \frac{n}{N}\right)^d. \quad \square$$

4. Fluid-Limit Proofs

In this section, we prove the fluid-limit results for the JSQ($d(N)$) scheme stated in Theorems 1 and 2. As mentioned in Subsection 2.4, the fluid limit for the ordinary JSQ policy is provided in Subsection 4.1, and in Subsection 4.2 we prove a universality result establishing that under the condition that $d(N) \rightarrow \infty$ as $N \rightarrow \infty$, the fluid limit for the JSQ($d(N)$) scheme coincides with that for the ordinary JSQ policy.

4.1. Fluid Limit of JSQ

In this section, we establish the fluid limit for the ordinary JSQ policy and the interchange of limits result stated in Proposition 1. In the proof we will leverage the timescale separation technique developed by Hunt and Kurtz (1994), suitably extended to an infinite-dimensional space. As mentioned in the introduction, to the best of our knowledge, this is the first time the transient fluid limit of the ordinary JSQ policy is rigorously established. We also observe that in order to exploit the coupling framework in Section 3.2 and, in particular, Proposition 3, we need convergence of tail sums. Thus, we need to establish the fluid convergence result with respect to the ℓ_1 topology, which makes the analysis technically challenging.

To leverage the timescale separation technique, note that the rate at which incoming tasks join a server with queue length i is determined only by the process $\mathbf{Z}^N(\cdot) = (Z_1^N(\cdot), \dots, Z_b^N(\cdot))$, where $Z_i^N(t) = N - Q_i^N(t)$, $i = 1, \dots, b$, represents the number of servers with queue lengths smaller than i at time t . Furthermore, the dynamics of the $\mathbf{Z}^N(\cdot)$ process can be described as

$$\mathbf{Z}^N \rightarrow \begin{cases} \mathbf{Z}^N + e_i & \text{at rate } N(q_i^N - q_{i+1}^N), \\ \mathbf{Z}^N - e_i & \text{at rate } N\lambda \mathbf{1}_{\{\mathbf{Z}^N \in \mathcal{R}_i\}}, \end{cases} \quad (17)$$

where e_i is the i th unit vector, and

$$\mathcal{R}_i := \{(z_1, z_2, \dots, z_b) : z_1 = \dots = z_{i-1} = 0 < z_i \leq z_{i+1} \leq \dots \leq z_b\} \in \mathcal{G}, \quad (18)$$

$i = 1, 2, \dots, b$, with the convention that Q_{b+1}^N is always taken to be zero, if $b < \infty$. Observe that in any time interval $[t, t + \varepsilon]$ of length $\varepsilon > 0$, the $\mathbf{Z}^N(\cdot)$ process experiences $O(\varepsilon N)$ events (arrivals and departures), whereas the $\mathbf{q}^N(\cdot)$ process can change by only $O(\varepsilon)$ amount. In other words, loosely speaking, around a “small” neighborhood of time t ,

the $q_i(t)$'s are constants, while as $N \rightarrow \infty$, the process $\mathbf{Z}^N(\cdot)$ behaves as a *timescaled* version of the following process:

$$\mathbf{Z}_{\mathbf{q}(t)} \rightarrow \begin{cases} \mathbf{Z}_{\mathbf{q}(t)} + e_i & \text{at rate } q_i(t) - q_{i+1}(t), \\ \mathbf{Z}_{\mathbf{q}(t)} - e_i & \text{at rate } \lambda \mathbf{1}_{[\mathbf{Z}_{\mathbf{q}(t)} \in \mathcal{R}_i]}. \end{cases} \quad (19)$$

Therefore, the $\mathbf{Z}^N(\cdot)$ process evolves on a much faster timescale than the $\mathbf{q}^N(\cdot)$ process. As a result, in the limit as $N \rightarrow \infty$, at each time point t , the $\mathbf{Z}^N(\cdot)$ process achieves stationarity depending on the instantaneous value of the $\mathbf{q}^N(\cdot)$ process, i.e., a separation of timescales takes place. In order to establish the timescale separation and the fluid limit results, we first write the evolution of the occupancy states in terms of a suitable random measure (see Equation (32)) and establish in Proposition 8 that the sequence of joint occupancy process and the random measure is relatively compact. We also characterize the limit of any convergence subsequence, where we invoke analogous arguments as used in the proofs of Hunt and Kurtz (1994, lemma 2) and Hunt and Kurtz (1994, theorem 3) to complete the proof of the separation of timescales. The proof of the fluid limit result is then completed by establishing uniqueness of the instantaneous stationary distribution achieved by the fast process, given any fluid-scaled occupancy state.

Denote by $\bar{\mathbb{Z}}_+$ the one-point compactification of the set of nonnegative integers \mathbb{Z}_+ , i.e., $\bar{\mathbb{Z}}_+ = \mathbb{Z}_+ \cup \{\infty\}$. Equip $\bar{\mathbb{Z}}_+$ with the order topology. Denote $G = \bar{\mathbb{Z}}_+^b$ equipped with the product topology, and with the Borel σ -algebra \mathcal{G} . Let us consider the G -valued process $\mathbf{Z}^N(s) := (Z_i^N(s))_{i \geq 1}$ as introduced above. Note that for the ordinary JSQ policy, the probability that a task arriving at (say) t_k is assigned to some server with queue length i is given by $p_{i-1}^N(\mathbf{Q}^N(t_k-)) = \mathbf{1}_{[\mathbf{Z}^N(t_k-) \in \mathcal{R}_i]}$, where \mathcal{R}_i is as in Equation (18). We prove the following fluid-limit result for the ordinary JSQ policy. Recall the definition of $m(\mathbf{q})$ in Subsection 2.2. If $m(\mathbf{q}) > 0$, then define

$$p_i(\mathbf{q}) = \begin{cases} \min \{ (1 - q_{m(\mathbf{q}+1)}) / \lambda, 1 \} & \text{for } i = m(\mathbf{q}) - 1, \\ 1 - p_{m(\mathbf{q})-1}(\mathbf{q}) & \text{for } i = m(\mathbf{q}), \\ 0 & \text{otherwise,} \end{cases} \quad (20)$$

and else, define $p_0(\mathbf{q}) = 1$ and $p_i(\mathbf{q}) = 0$ for all $i = 1, \dots, b$.

Theorem 5 (Fluid Limit of JSQ). *Assume $\mathbf{q}^N(0) \xrightarrow{\mathbb{P}} \mathbf{q}^\infty$ in \mathcal{S} and $\lambda(N)/N \rightarrow \lambda > 0$ as $N \rightarrow \infty$. Then, any subsequence of the sequence of processes $\{\mathbf{q}^N(t)\}_{t \geq 0}$ for the ordinary JSQ policy has a further subsequence that converges weakly with respect to the Skorohod J_1 topology to the limit $\{\mathbf{q}(t)\}_{t \geq 0}$ satisfying the following system of integral equations:*

$$q_i(t) = q_i(0) + \lambda \int_0^t p_{i-1}(\mathbf{q}(s)) ds - \int_0^t (q_i(s) - q_{i+1}(s)) ds, \quad i = 1, 2, \dots, b, \quad (21)$$

where $\mathbf{q}(0) = \mathbf{q}^\infty$ and the coefficients $p_i(\cdot)$ are as defined in Equation (20).

The rest of this section will be devoted in the proof of Theorem 5. First, we construct the martingale representation of the occupancy state process $\mathbf{Q}^N(\cdot)$. Note that the component $Q_i^N(t)$ satisfies the identity relation

$$Q_i^N(t) = Q_i^N(0) + A_i^N(t) - D_i^N(t), \quad \text{for } i = 1, \dots, b, \quad (22)$$

where

$$\begin{aligned} A_i^N(t) &= \text{number of arrivals during } [0, t] \text{ to some server with queue length } i - 1, \\ D_i^N(t) &= \text{number of departures during } [0, t] \text{ from some server with queue length } i. \end{aligned}$$

We can express $A_i^N(t)$ and $D_i^N(t)$ as

$$\begin{aligned} A_i^N(t) &= \mathcal{N}_{A,i} \left(\lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds \right), \\ D_i^N(t) &= \mathcal{N}_{D,i} \left(\int_0^t (Q_i^N(s) - Q_{i+1}^N(s)) ds \right), \end{aligned}$$

where $\mathcal{N}_{A,i}$ and $\mathcal{N}_{D,i}$ are mutually independent unit-rate Poisson processes, $i = 1, 2, \dots, b$. Define the following sigma fields:

$$\begin{aligned} \mathcal{A}_i^N(t) &:= \sigma(A_i^N(s) : 0 \leq s \leq t), \\ \mathcal{D}_i^N(t) &:= \sigma(D_i^N(s) : 0 \leq s \leq t), \text{ for } i = 1, \dots, b, \end{aligned}$$

and the filtration $\mathbf{F}^N \equiv \{\mathcal{F}_t^N : t \geq 0\}$ with

$$\mathcal{F}_t^N := \bigvee_{i=1}^{\infty} [\mathcal{A}_i^N(t) \vee \mathcal{D}_i^N(t)] \quad (23)$$

augmented by all the null sets. Now, we have the following martingale decomposition from the random time change of a unit-rate Poisson process result of Pang et al. (2007, lemma 3.2).

Proposition 6 (Martingale Decomposition). *The following are \mathbf{F}^N martingales, for $i \geq 1$:*

$$\begin{aligned} M_{A,i}^N(t) &:= \mathcal{N}_{A,i} \left(\lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds \right) - \lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds, \\ M_{D,i}^N(t) &:= \mathcal{N}_{D,i} \left(\int_0^t (Q_i^N(s) - Q_{i+1}^N(s)) ds \right) - \int_0^t (Q_i^N(s) - Q_{i+1}^N(s)) ds, \end{aligned} \quad (24)$$

with respective compensator and predictable quadratic variation processes given by

$$\begin{aligned} \langle M_{A,i}^N \rangle(t) &:= \lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s-)) ds, \\ \langle M_{D,i}^N \rangle(t) &:= \int_0^t (Q_i^N(s) - Q_{i+1}^N(s)) ds. \end{aligned}$$

Therefore, finally, we have the following martingale representation of the N th process:

$$\begin{aligned} Q_i^N(t) &= Q_i^N(0) + \lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds - \int_0^t (Q_i^N(s) - Q_{i+1}^N(s)) ds \\ &\quad + (M_{A,i}^N(t) - M_{D,i}^N(t)), \quad t \geq 0, \quad i = 1, \dots, b. \end{aligned} \quad (25)$$

In the proposition below, we prove that the martingale part vanishes in ℓ_1 when scaled by N .

Proposition 7 (Convergence of Martingales).

$$\left\{ \frac{1}{N} \sum_{i \geq 1} (|M_{A,i}^N(t)| + |M_{D,i}^N(t)|) \right\}_{t \geq 0} \xrightarrow{d} \{m(t)\}_{t \geq 0} \equiv 0.$$

Proof. The proof follows using the same line of arguments as in the proof of Mitzenmacher (1996, theorem 3.13) and hence is sketched only briefly for the sake of completeness. Fix any $T \geq 0$, and observe that

$$\lim_{N \rightarrow \infty} \sup_{t \in [0, T]} \frac{1}{N} \sum_{i \geq 1} |M_{A,i}^N(t)| \quad (26)$$

$$= \lim_{N \rightarrow \infty} \sup_{t \in [0, T]} \frac{1}{N} \left(\sum_{i \geq 1} \left| \mathcal{N}_{A,i} \left(\lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds \right) - \lambda(N) \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds \right| \right) \quad (27)$$

$$\leq \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i \geq 1} \mathcal{N}_{A,i} \left(\lambda(N) \int_0^T p_{i-1}^N(\mathbf{Q}^N(s)) ds \right) + \lambda T. \quad (28)$$

Because $N^{-1} \sum_{i \geq 1} \lambda(N) \int_0^T p_{i-1}^N(\mathbf{Q}^N(s)) ds \rightarrow \lambda t < \infty$, the $\lim_{N \rightarrow \infty}$ and $\sum_{i \geq 1}$ above can be interchanged in Equation (28), and hence in Equation (26). Now, for each $i \geq 1$, from Doob's inequality (Liptser and Shiryaev 1989, theorem 1.9.1.3), we have for any $\epsilon > 0$,

$$\begin{aligned} \mathbb{P} \left(\sup_{t \in [0, T]} \frac{1}{N} M_{A,i}^N(t) \geq \epsilon \right) &= \mathbb{P} \left(\sup_{t \in [0, T]} M_{A,i}^N(t) \geq N\epsilon \right) \leq \frac{1}{N^2 \epsilon^2} \mathbb{E}(\langle M_{A,i}^N \rangle(T)) \\ &\leq \frac{1}{N \epsilon^2} \int_0^T p_{i-1}(\mathbf{Q}^N(s-)) \lambda(N) ds \leq \frac{\lambda T}{N \epsilon^2} \rightarrow 0, \text{ as } N \rightarrow \infty. \end{aligned}$$

Thus, $\sup_{t \in [0, T]} N^{-1} M_{A,i}^N(t) \xrightarrow{\mathbb{P}} 0$, and hence, $\sup_{t \in [0, T]} N^{-1} \sum_{i \geq 1} |M_{A,i}^N(t)| \xrightarrow{\mathbb{P}} 0$. Using similar arguments as above, we can also show that $\sup_{t \in [0, T]} N^{-1} \sum_{i \geq 1} |M_{D,i}^N(t)| \xrightarrow{\mathbb{P}} 0$, and the proof is complete. \square

Now, we prove the relative compactness of the sequence of fluid-scaled processes. Recall that we denote all the fluid-scaled quantities by their respective small letters, e.g., $\mathbf{q}^N(t) := \mathbf{Q}^N(t)/N$, componentwise, i.e., $q_i^N(t) := Q_i^N(t)/N$ for $i \geq 1$. Therefore, the martingale representation in Equation (25), can be written as

$$\begin{aligned} q_i^N(t) &= q_i^N(0) + \frac{\lambda(N)}{N} \int_0^t p_{i-1}^N(\mathbf{Q}^N(s)) ds - \int_0^t (q_i^N(s) - q_{i+1}^N(s)) ds \\ &\quad + \frac{1}{N} (M_{A,i}^N(t) - M_{D,i}^N(t)), \quad i = 1, 2, \dots, b, \end{aligned} \quad (29)$$

or equivalently,

$$\begin{aligned} q_i^N(t) &= q_i^N(0) + \frac{\lambda(N)}{N} \int_0^t \mathbf{1}_{[Z^N(s) \in \mathcal{R}_i]} ds - \int_0^t (q_i^N(s) - q_{i+1}^N(s)) ds \\ &\quad + \frac{1}{N} (M_{A,i}^N(t) - M_{D,i}^N(t)), \quad i = 1, 2, \dots, b. \end{aligned} \quad (30)$$

Now, we consider the Markov process $(\mathbf{q}^N, \mathbf{Z}^N)(\cdot)$ defined on $\mathcal{S} \times G$. Define a random measure α^N on the measurable space $([0, \infty) \times G, \mathcal{C} \otimes \mathcal{G})$, when $[0, \infty)$ is endowed with Borel sigma algebra \mathcal{C} by

$$\alpha^N(A_1 \times A_2) := \int_{A_1} \mathbf{1}_{[Z^N(s) \in A_2]} ds, \quad (31)$$

for $A_1 \in \mathcal{C}$ and $A_2 \in \mathcal{G}$. Then, the representation in Equation (30) can be written in terms of the random measure as

$$\begin{aligned} q_i^N(t) &= q_i^N(0) + \lambda \alpha^N([0, t] \times \mathcal{R}_i) - \int_0^t (q_i^N(s) - q_{i+1}^N(s)) ds \\ &\quad + \frac{1}{N} (M_{A,i}^N(t) - M_{D,i}^N(t)), \quad i = 1, 2, \dots, b. \end{aligned} \quad (32)$$

Let \mathcal{L} denote the space of all measures on $[0, \infty) \times G$ satisfying $\gamma([0, t], G) = t$, endowed with the topology corresponding to weak convergence of measures restricted to $[0, t] \times G$ for each t .

Proposition 8 (Relative Compactness). *Assume $\mathbf{q}^N(0) \xrightarrow{d} \mathbf{q}^\infty \in \mathcal{S}$ as $N \rightarrow \infty$, then $\{(\mathbf{q}^N(\cdot), \alpha^N)\}_{N \geq 1}$ is a relatively compact sequence in $D_{\mathcal{G}}[0, \infty) \times \mathcal{L}$ and the limit $(\mathbf{q}(\cdot), \alpha)$ of any convergent subsequence satisfies*

$$q_i(t) = q_i^\infty + \lambda \alpha([0, t] \times \mathcal{R}_i) - \int_0^t (q_i(s) - q_{i+1}(s)) ds, \quad i = 1, 2, \dots, b. \quad (33)$$

To prove Proposition 8, we will verify the relative compactness conditions from Ethier and Kurtz (2009). Let (E, r) be a complete and separable metric space. For any $x \in D_E[0, \infty)$, $\delta > 0$ and $T > 0$, define

$$w'(x, \delta, T) = \inf_{\{t_i\}} \max_i \sup_{s, t \in [t_{i-1}, t_i]} r(x(s), x(t)), \quad (34)$$

where $\{t_i\}$ ranges over all partitions of the form $0 = t_0 < t_1 < \dots < t_{n-1} < T \leq t_n$ with $\min_{1 \leq i \leq n} (t_i - t_{i-1}) > \delta$ and $n \geq 1$. Below we state the conditions for the sake of completeness.

Theorem 6 (Ethier and Kurtz 2009, corollary 3.7.4). *Let (E, r) be complete and separable, and let $\{X_n\}_{n \geq 1}$ be a family of processes with sample paths in $D_E[0, \infty)$. Then, $\{X_n\}_{n \geq 1}$ is relatively compact if and only if the following two conditions hold:*

(a) *For every $\eta > 0$ and rational $t \geq 0$, there exists a compact set $\Gamma_{\eta, t} \subset E$ such that*

$$\liminf_{n \rightarrow \infty} \mathbb{P}(X_n(t) \in \Gamma_{\eta, t}) \geq 1 - \eta.$$

(b) *For every $\eta > 0$ and $T > 0$, there exists $\delta > 0$ such that*

$$\overline{\lim}_{n \rightarrow \infty} \mathbb{P}(w'(X_n, \delta, T) \geq \eta) \leq \eta.$$

In order to prove the relative compactness, we will need the next three lemmas: Lemma 1 characterizes the relatively compact subsets of \mathcal{S} , Lemma 2 provides a necessary and sufficient criterion for a sequence of ℓ_1 -valued random variables to be tight, and Lemma 3 is needed to ensure that at all finite times t , the occupancy state process lies in some compact set (possibly depending upon t).

Lemma 1 (Compact Subsets of \mathcal{S}). *Assume $b = \infty$. A set $K \subseteq \mathcal{S}$ is relatively compact in \mathcal{S} with respect to the ℓ_1 topology if and only if*

$$\limsup_{k \rightarrow \infty} \sum_{\mathbf{x} \in K} \sum_{i=k}^{\infty} x_i = 0. \quad (35)$$

Proof. For the if part, fix any $K \subseteq \mathcal{S}$ satisfying Equation (35). We will show that any sequence $\{\mathbf{x}^n\}_{n \geq 1}$ in K has a Cauchy subsequence. Because the ℓ_1 space is complete, this will then imply that $\{\mathbf{x}^n\}_{n \geq 1}$ has a convergent subsequence with the limit in \overline{K} , which will complete the proof.

To show the existence of a Cauchy sequence, fix any $\varepsilon > 0$ and choose $k \geq 1$ (depending on ε) such that

$$\sum_{i \geq k} |x_i^n| < \frac{\varepsilon}{4} \quad \forall n \geq 1. \quad (36)$$

Now, observe that the set of first coordinates $\{x_1^n\}_{n \geq 1}$ is a sequence in $[0, 1]$ and hence has a convergent subsequence. Along that subsequence, the set of the second coordinates has a further convergent subsequence. Proceeding this way, we can get a subsequence along which the first $k - 1$ coordinates converge. Therefore, depending upon ε , an $N' \in \mathbb{N}$ can be chosen, such that

$$\sum_{i < k} |x_i^n - x_i^m| < \frac{\varepsilon}{2} \quad \forall m, n \geq N'. \quad (37)$$

Therefore, Equations (36) and (37) yield for all $n \geq \max\{N, N'\}$,

$$\begin{aligned} \|\mathbf{x}^n - \mathbf{x}^m\|_1 &= \sum_{i \geq 1} |x_i^n - x_i^m| \leq \sum_{i < k} |x_i^n - x_i^m| + \sum_{i \geq k} |x_i^n - x_i^m| \\ &\leq \sum_{i < k} |x_i^n - x_i^m| + \sum_{i \geq k} x_i^n + \sum_{i \geq k} x_i^m < \varepsilon \end{aligned}$$

along the above suitably constructed subsequence. Now that the limit point is in \mathcal{S} follows from the completeness of the ℓ_1 space and the fact that \mathcal{S} is a closed subset of ℓ_1 . Indeed, because the ℓ_1 topology is finer than the product topology, any set that is closed with respect to the product topology is closed with respect to the ℓ_1 topology, and observe that \mathcal{S} is closed with respect to the product topology.

For the only if part, let $K \subseteq \mathcal{S}$ be relatively compact, and on the contrary, assume that there exists an $\varepsilon > 0$, such that

$$\limsup_{k \rightarrow \infty} \sum_{\mathbf{x} \in K} \sum_{i=k}^{\infty} x_i \geq \varepsilon. \quad (38)$$

Therefore, for each $k \geq 1$, there exists $\mathbf{x}^{(k)} \in K$, such that $\sum_{i=k}^{\infty} x_i^{(k)} \geq \varepsilon/2$. Consider any limit point \mathbf{x}^* of the sequence $\{\mathbf{x}^{(k)}\}_{k \geq 1}$, and note that $\sum_{i=j}^{\infty} x_i^* \geq \varepsilon/2$ for all $j \geq 1$. This contradicts that $\mathbf{x}^* \in \ell_1$, and the proof is complete. \square

Lemma 2 (Criterion for ℓ_1 Tightness). *Let $\{\mathbf{X}^N\}_{N \geq 1}$ be a sequence of random variables in \mathcal{S} . Then, the following are equivalent:*

(i) $\{\mathbf{X}^N\}_{N \geq 1}$ is tight with respect to the product topology, and for all $\varepsilon > 0$,

$$\lim_{k \rightarrow \infty} \overline{\lim}_{N \rightarrow \infty} \mathbb{P}\left(\sum_{i \geq k} X_i^N > \varepsilon\right) = 0. \quad (39)$$

(ii) $\{\mathbf{X}^N\}_{N \geq 1}$ is tight with respect to the ℓ_1 topology.

Proof. To prove (i) \implies (ii), for any $\varepsilon > 0$, we will construct a relatively compact set compact set $K(\varepsilon)$ such that

$$\mathbb{P}\left(\mathbf{X}^N \notin \overline{K(\varepsilon)}\right) < \varepsilon \quad \text{for all } N.$$

Observe from Equation (39) that for all $\varepsilon > 0$, there exists an $r(\varepsilon) \geq 1$, such that

$$\overline{\lim}_{N \rightarrow \infty} \mathbb{P} \left(\sum_{i \geq r(\varepsilon)} X_i^N > \varepsilon \right) < \varepsilon,$$

and with it an $N(\varepsilon) \geq 1$, such that

$$\mathbb{P} \left(\sum_{i \geq k(\varepsilon)} X_i^N > \varepsilon \right) < \varepsilon \quad \text{for all } N > N(\varepsilon).$$

Furthermore, because $\{\mathbf{X}^1, \mathbf{X}^2, \dots, \mathbf{X}^{N(\varepsilon)}\}$ is a finite set of ℓ_1 -valued random variables, there exists $k(\varepsilon) \geq r(\varepsilon)$, such that

$$\mathbb{P} \left(\sum_{i \geq k(\varepsilon)} X_i^N > \varepsilon \right) < \varepsilon \quad \text{for all } N.$$

Thus, there exists an increasing sequence $\{k(n)\}_{n \geq 1}$ such that

$$\mathbb{P} \left(\sum_{i \geq k(n)} X_i^N > \frac{\varepsilon}{2^n} \right) < \frac{\varepsilon}{2^n} \quad \text{for all } N.$$

Define the set $K(\varepsilon)$ as

$$K(\varepsilon) := \left\{ \mathbf{x} \in \mathcal{S} : \sum_{i \geq k(n)} x_i \leq \frac{\varepsilon}{2^n} \quad \text{for all } n \geq 1 \right\}.$$

Because of Lemma 1, we know $K(\varepsilon)$ is relatively compact in ℓ_1 . Also,

$$\mathbb{P}(\mathbf{X}^N \notin \overline{K(\varepsilon)}) = \mathbb{P} \left(\bigcup_{n \geq 1} \left\{ \sum_{i \geq k(n)} X_i^N > \frac{\varepsilon}{2^n} \right\} \right) \leq \sum_{n \geq 1} \mathbb{P} \left(\sum_{i \geq k(n)} X_i^N > \frac{\varepsilon}{2^n} \right) < \varepsilon.$$

To prove (ii) \implies (i), first observe that a sequence of random variables is tight with respect to the ℓ_1 topology implies that it must be tight with respect to the product topology. Now, assume on the contrary to Equation (39) that there exists $\varepsilon > 0$, such that

$$\lim_{k \rightarrow \infty} \overline{\lim}_{N \rightarrow \infty} \mathbb{P} \left(\sum_{i \geq k} X_i^N > \varepsilon \right) > \varepsilon. \quad (40)$$

Because $\{\mathbf{X}^N\}_{N \geq 1}$ is tight with respect to ℓ_1 topology, take any convergent subsequence $\{\mathbf{X}^{N(n)}\}_{n \geq 1}$ with \mathbf{X}^* being a random variable following the limiting measure. In that case, observe that Equation (40) implies $\mathbb{P}(\sum_{i \geq k} X_i^* > \varepsilon/2) > \varepsilon$ for all $k \geq 1$, which leads to a contradiction because \mathbf{X}^* is an ℓ_1 -valued random variable. \square

Lemma 3. For any $\mathbf{q} \in \mathcal{S}$, assume that $\mathbf{q}^N(0) \xrightarrow{d} \mathbf{q}^\infty$, as $N \rightarrow \infty$. Then, for any $t \geq 0$, there exists $M(t, \mathbf{q}^\infty) \geq 1$, such that under the JSQ policy, with probability tending to one as $N \rightarrow \infty$, no arriving task is assigned to a server with queue length $M(t, \mathbf{q}^\infty) - 1$ up to time t .

Proof. Let $A^N(t)$ be the cumulative number of tasks arriving up to time t . Because the arrival rate is $\lambda(N)$ and $\lambda(N)/N \rightarrow \lambda$, as $N \rightarrow \infty$, for any $\varepsilon > 0$,

$$\mathbb{P}(A^N(t) \geq (\lambda t + \varepsilon)N) \rightarrow 0 \quad \text{as } N \rightarrow \infty.$$

Define $M(t, \mathbf{q}^\infty) := \min \{k \geq 1 : \sum_{i=1}^{k-1} (1 - q_i^\infty) > \lambda t\}$, and choose

$$\varepsilon = \sum_{i=1}^{M(t, \mathbf{q}^\infty)-1} (1 - q_i^\infty) - \lambda t > 0.$$

Note that because $\mathbf{q}^\infty \in \mathcal{S} \subset \ell_1$, $M(t, \mathbf{q}^\infty)$ exists and is finite for all $t \geq 0$. We now claim that the probability that in the interval $[0, t]$ a task is assigned to some server with queue length $M(t, \mathbf{q}^\infty)$ tends to 0, as $N \rightarrow \infty$. Indeed, in order for a task to be assigned to some server with queue length $M(t, \mathbf{q}^\infty) - 1$, all the servers must have queue length

$M(t, \mathbf{q}^\infty) - 1$ or larger. Now, the minimum number of tasks required for this, is given by $\sum_{i=1}^{M(t, \mathbf{q}^\infty)-1} (N - Q_i^N(0))$. Therefore, the proof is complete by observing that

$$\mathbb{P}\left(A^N(t) \geq \sum_{i=1}^{M(t, \mathbf{q}^\infty)-1} (N - Q_i^N(0))\right) = \mathbb{P}\left(A^N(t) \geq \left(\lambda t + \frac{\varepsilon}{2}\right)N\right) \rightarrow 0, \quad \text{as } N \rightarrow \infty. \quad \square$$

Proof of Proposition 8. The proof goes in two steps. We first prove the relative compactness and then show that the limit satisfies Equation (33).

Observe from Ethier and Kurtz (2009, proposition 3.2.4) that, to prove the relative compactness of the sequence of processes $\{(\mathbf{q}^N(\cdot), \alpha^N)\}_{N \geq 1}$, it is enough to prove relative compactness of the individual components. Note that, from Prohorov’s theorem (Ethier and Kurtz 2009, theorem 3.2.2), \mathfrak{L} is compact because G is compact. Now, relative compactness of $\{\alpha^N\}_{N \geq 1}$ follows from the compactness of \mathfrak{L} under the topology of weak convergence of measures and Prohorov’s theorem. To claim the relative compactness of $\{\mathbf{q}^N(\cdot)\}_{N \geq 1}$, we will verify the conditions of Theorem 6.

Observe that Theorem 6(a) requires to show tightness of the sequence $\{\mathbf{q}^N(t)\}_{N \geq 1}$ for each fixed (rational) $t \geq 0$. Fix any $t \geq 0$. Because of Lemma 3, we know

$$\lim_{N \rightarrow \infty} \mathbb{P}(q_i^N(t) \leq q_i^N(0), \quad \forall i \geq M(t, \mathbf{q}^\infty)) = 1.$$

Also, $\mathbf{q}^N(0) \xrightarrow{d} \mathbf{q}^\infty$ with respect to the ℓ_1 topology. In particular, $\{\mathbf{q}^N(0)\}_{N \geq 1}$ is tight in ℓ_1 . Therefore, using (ii) \implies (i) in Lemma 2 we obtain, for any $\varepsilon > 0$,

$$\lim_{k \rightarrow \infty} \overline{\lim}_{N \rightarrow \infty} \mathbb{P}\left(\sum_{i \geq k} q_i^N(t) > \varepsilon\right) \leq \lim_{k \rightarrow \infty} \overline{\lim}_{N \rightarrow \infty} \mathbb{P}\left(\sum_{i \geq k} q_i^N(0) > \varepsilon\right) = 0.$$

Also, because $\mathbf{q}^N(t) \in \mathcal{G} \subseteq [0, 1]^b$, which is compact with respect to the product topology, $\{\mathbf{q}^N(t)\}_{N \geq 1}$ is tight with respect to the product topology. Hence, using (i) \implies (ii) in Lemma 2, we conclude that the sequence $\{\mathbf{q}^N(t)\}_{N \geq 1}$ is tight in ℓ_1 . For condition (b), first note that for all $i = 1, \dots, b$,

$$\begin{aligned} |q_i^N(t_1) - q_i^N(t_2)| &\leq \lambda \alpha^N([t_1, t_2] \times \mathcal{R}_i) + \int_{t_1}^{t_2} (q_i^N(s) - q_{i+1}^N(s)) ds \\ &\quad + \frac{1}{N} |M_{A,i}^N(t_1) - M_{D,i}^N(t_1) - M_{A,i}^N(t_2) + M_{D,i}^N(t_2)| + o(1). \end{aligned}$$

Thus,

$$\begin{aligned} &\|\mathbf{q}^N(t_1) - \mathbf{q}^N(t_2)\|_1 \\ &\leq \lambda \sum_{i=1}^b \alpha^N([t_1, t_2] \times \mathcal{R}_i) + \int_{t_1}^{t_2} \sum_{i=1}^b (q_i^N(s) - q_{i+1}^N(s)) ds \\ &\quad + \frac{1}{N} \sum_{i=1}^b |M_{A,i}^N(t_1) - M_{D,i}^N(t_1) - M_{A,i}^N(t_2) + M_{D,i}^N(t_2)| + o(1) \\ &\leq \lambda(t_1 - t_2) + \int_{t_1}^{t_2} q_1^N(s) ds + \frac{1}{N} \sum_{i=1}^b |M_{A,i}^N(t_1) - M_{D,i}^N(t_1) - M_{A,i}^N(t_2) + M_{D,i}^N(t_2)| + o(1) \\ &\leq (\lambda + 1)(t_1 - t_2) + \frac{1}{N} \sum_{i=1}^b |M_{A,i}^N(t_1) - M_{D,i}^N(t_1) - M_{A,i}^N(t_2) + M_{D,i}^N(t_2)| + o(1). \end{aligned} \tag{41}$$

Now, from the ℓ_1 convergence of scaled martingales in Proposition 7, we get, for any $T \geq 0$,

$$\sup_{t \in [0, T]} \frac{1}{N} \sum_{i=1}^b |M_{A,i}^N(t_1) - M_{D,i}^N(t_1) - M_{A,i}^N(t_2) + M_{D,i}^N(t_2)| \xrightarrow{\mathbb{P}} 0.$$

Observe that the proof of the relative compactness of $\{\mathbf{q}^N(t)\}_{t \geq 0}$ is complete if we show that for any $\eta > 0$, there exists a $\delta > 0$ and a finite partition $(t_j)_{j=1}^n$ of $[0, T]$ with $\min_j |t_j - t_{j-1}| > \delta$ such that

$$\overline{\lim}_{N \rightarrow \infty} \mathbb{P} \left(\max_j \sup_{s, t \in [t_{j-1}, t_j]} \|\mathbf{q}^N(s) - \mathbf{q}^N(t)\|_1 \geq \eta \right) < \eta. \tag{42}$$

Now, Equation (41) implies that, for any finite partition $(t_j)_{j=1}^n$ of $[0, T]$,

$$\max_j \sup_{s, t \in [t_{j-1}, t_j]} \|\mathbf{q}^N(s) - \mathbf{q}^N(t)\|_1 \leq (\lambda + 1) \max_j (t_j - t_{j-1}) + \zeta_N,$$

where $\mathbb{P}(\zeta_N > \eta/2) < \eta$ for all sufficiently large N . Now, take $\delta = \eta/(4(\lambda + 1))$ and any partition with $\max_j (t_j - t_{j-1}) < \eta/(2(\lambda + 1))$ and $\min_j (t_j - t_{j-1}) > \delta$. Now, on the event $\{\zeta_N \leq \eta/2\}$,

$$\max_i \sup_{s, t \in [t_{i-1}, t_i]} \|\mathbf{q}^N(s) - \mathbf{q}^N(t)\|_1 \leq \eta.$$

Therefore, for all sufficiently large N ,

$$\mathbb{P} \left(\max_j \sup_{s, t \in [t_{j-1}, t_j]} \|\mathbf{q}^N(s) - \mathbf{q}^N(t)\|_1 \geq \eta \right) \leq \mathbb{P}(\zeta_N > \eta/2) \leq \eta.$$

To prove that the limit $(\mathbf{q}(\cdot), \alpha)$ of any convergent subsequence satisfies Equation (33), we will use the continuous-mapping theorem (Whitt 2002, theorem 3.4.1). Specifically, we will show that the right side of Equation (32) is a continuous map of suitable arguments. Let $\{\mathbf{q}(t)\}_{t \geq 0}$ and $\{\mathbf{y}(t)\}_{t \geq 0}$ be an \mathcal{S} -valued and an ℓ_1 -valued càdlàg function, respectively. Also, let α be a measure on the measurable space $([0, \infty) \times G, \mathcal{C} \otimes \mathcal{G})$. Then, for $\mathbf{q}^0 \in \mathcal{S}$, define for $i \geq 1$,

$$F_i(\mathbf{q}, \alpha, \mathbf{q}^0, \mathbf{y})(t) := q_i^0 + y_i(t) + \lambda \alpha([0, t] \times \mathcal{R}_i) - \int_0^t (q_i(s) - q_{i+1}(s)) ds.$$

Observe that it is enough to show $\mathbf{F} = (F_1, \dots, F_b)$ is a continuous operator. Indeed, in that case the right side of Equation (32) can be written as $\mathbf{F}(\mathbf{q}^N, \alpha^N, \mathbf{q}^N(0), \mathbf{y}^N)$, where $\mathbf{y}^N = (y_1^N, \dots, y_b^N)$ with $y_i^N = (M_{A,i}^N - M_{D,i}^N)/N$, and because each argument converges, we will get the convergence to the right side of Equation (33). Therefore, we now prove the continuity of \mathbf{F} below. In particular assume that (a) the sequence of processes $\{(\mathbf{q}^N, \mathbf{y}^N)\}_{N \geq 1}$ converges to (\mathbf{q}, \mathbf{y}) with respect to the ℓ_1 topology, (b) for any fixed $t \geq 0$, the sequence $\{(\alpha^N([0, t], \mathcal{R}_i))_{i \geq 1}\}_{N \geq 1}$ in ℓ_1 converges to $(\alpha([0, t], \mathcal{R}_i))_{i \geq 1}$, and (c) the sequence of \mathcal{S} -valued random variables $\mathbf{q}^N(0)$ converges to $\mathbf{q}(0)$ with respect to the ℓ_1 topology.

Fix any $T \geq 0$ and $\varepsilon > 0$.

(i) Because of (a) above, choose $N_1 \in \mathbb{N}$, such that $\sup_{t \in [0, T]} \|\mathbf{q}^N(t) - \mathbf{q}(t)\|_1 < \varepsilon/(4T)$. In that case, observe that

$$\sup_{t \in [0, T]} \int_0^t |q_1^N(t) - q_1(t)| ds \leq T \sup_{t \in [0, T]} \|\mathbf{q}^N(t) - \mathbf{q}(t)\|_1 < \frac{\varepsilon}{4}.$$

(ii) Again, because of (a), choose $N_2 \in \mathbb{N}$, such that $\sup_{t \in [0, T]} \|\mathbf{y}^N(t) - \mathbf{y}(t)\|_1 < \varepsilon/4$.

(iii) We now claim that for the $\varepsilon > 0$ given above there is an $N_3 \in \mathbb{N}$ such that

$$\lambda \sum_{i \geq 1} |\alpha^N([0, T] \times \mathcal{R}_i) - \alpha([0, T] \times \mathcal{R}_i)| < \frac{\varepsilon}{4}. \tag{43}$$

Observe that we only know the weak convergence of the sequence of measures α^N and therefore we cannot directly make assumption (b) above. We therefore now proceed to show that assumption (b) is valid in our case and that it follows from weak convergence. Indeed, because $\mathbf{q}^\infty \in \mathcal{S} \subseteq \ell_1$, there exists $\hat{M}(\mathbf{q}^\infty)$, such that $q_{\hat{M}(\mathbf{q}^\infty)}^\infty < 1$, and consequently $q_i^\infty < 1$ for all $i \geq \hat{M}(\mathbf{q}^\infty)$. Also, because of Lemma 3,

$$\lim_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} q_i^N(t) \leq q_i^N(0) \quad \text{for all } i \geq \hat{M}(T, \mathbf{q}^\infty) \right) = 1.$$

Thus, if $N_0 := \max\{\hat{M}(\mathbf{q}^\infty), M(T, \mathbf{q}^\infty)\}$, then

$$\lim_{N \rightarrow \infty} \mathbb{P}\left(\sup_{t \in [0, T]} q_i^N(t) < 1 \quad \text{for all } i \geq N_0\right) = 1.$$

This implies

$$\sum_{i \geq N_0} \alpha^N([0, T] \times \mathcal{R}_i) \xrightarrow{\mathbb{P}} \sum_{i \geq N_0} \alpha([0, T] \times \mathcal{R}_i) = 0.$$

Also, because of weak convergence of α^N ,

$$\sum_{i < N_0} \alpha^N([0, T] \times \mathcal{R}_i) \xrightarrow{\mathbb{P}} \sum_{i < N_0} \alpha([0, T] \times \mathcal{R}_i).$$

(iv) Finally, because of (c), choose $N_4 \in \mathbb{N}$, such that $\|\mathbf{q}^N(0) - \mathbf{q}(0)\|_1 < \varepsilon/4$. Let $\hat{N} = \max\{N_1, N_2, N_3, N_4\}$, then for $N \geq \hat{N}$,

$$\sup_{t \in [0, T]} \|\mathbf{F}(\mathbf{q}^N, \alpha^N, \mathbf{q}^N(0), \mathbf{y}^N) - \mathbf{F}(\mathbf{q}, \alpha, \mathbf{q}(0), \mathbf{y})\|_1(t) < \varepsilon.$$

Thus, the proof of continuity of \mathbf{F} is complete. \square

To characterize the limit in Equation (33), for any $\mathbf{q} \in \mathcal{S}$, define the Markov process $\mathbf{Z}_{\mathbf{q}}$ on G as

$$\mathbf{Z}_{\mathbf{q}} \rightarrow \begin{cases} \mathbf{Z}_{\mathbf{q}} + e_i & \text{at rate } q_i - q_{i+1}, \\ \mathbf{Z}_{\mathbf{q}} - e_i & \text{at rate } \lambda \mathbf{1}_{[\mathbf{Z}_{\mathbf{q}} \in \mathcal{R}_i]}, \end{cases} \quad (44)$$

where e_i is the i th unit vector, $i = 1, \dots, b$.

Proof of Theorem 5. Having proved the relative compactness in Proposition 8, it follows from analogous arguments as used in the proofs of Hunt and Kurtz (1994, lemma 2) and Hunt and Kurtz (1994, theorem 3), that the limit of any convergent subsequence of the sequence of processes $\{\mathbf{q}^N(t)\}_{t \geq 0}$ satisfies

$$q_i(t) = q_i(0) + \lambda \int_0^t \pi_{\mathbf{q}(s)}(\mathcal{R}_i) ds - \int_0^t (q_i(s) - q_{i+1}(s)) ds, \quad i = 1, 2, \dots, b, \quad (45)$$

for some stationary measure $\pi_{\mathbf{q}(t)}$ of the Markov process $\mathbf{Z}_{\mathbf{q}(t)}$ described in Equation (44) satisfying $\pi_{\mathbf{q}}\{\mathbf{Z} : Z_i = \infty\} = 1$ if $q_i < 1$.

Now, it remains to show that $\mathbf{q}(t)$ uniquely determines $\pi_{\mathbf{q}(t)}$ and that $\pi_{\mathbf{q}(s)}(\mathcal{R}_i) = p_{i-1}(\mathbf{q}(s))$ described in Equation (20). As mentioned earlier, in this proof we will now assume the specific assignment probabilities in Equation (18), corresponding to the ordinary JSQ policy. To see this, fix any $\mathbf{q} = (q_1, \dots, q_b) \in \mathcal{S}$. Observe that because of summability of the components of \mathbf{q} , there exists $0 \leq m < \infty$, such that $q_{m+1} < 1$ and $q_1 = \dots = q_m = 1$, with the convention that $q_0 \equiv 1$ and $q_{b+1} \equiv 0$ if $b < \infty$. In that case,

$$\pi_{\mathbf{q}}(\{Z_{m+1} = \infty, Z_{m+2} = \infty, \dots, Z_b = \infty\}) = 1. \quad (46)$$

Thus, if $q_1 < 1$, then Equation (46) implies $\pi_{\mathbf{q}}(\mathcal{R}_1) = 1$ and $\pi_{\mathbf{q}}(\mathcal{R}_i) = 0$ for all $i \geq 2$.

Now, assume $m \geq 1$. Note that $q_i = 1$ forces $dq_i/dt \leq 0$, i.e., $\lambda \pi_{\mathbf{q}}(\mathcal{R}_i) \leq q_i - q_{i+1}$ for all $i = 1, \dots, m$, and, in particular, $\pi_{\mathbf{q}}(\mathcal{R}_i) = 0$ for all $i = 1, \dots, m-1$. Thus,

$$\pi_{\mathbf{q}}(\{Z_1 = 0, Z_2 = 0, \dots, Z_{m-1} = 0\}) = 1.$$

Therefore, $\pi_{\mathbf{q}}$ is determined only by the stationary distribution of the m th component, which can be described as a birth-death process

$$Z \rightarrow \begin{cases} Z + 1 & \text{at rate } q_m - q_{m+1} \\ Z - 1 & \text{at rate } \lambda \mathbf{1}_{[Z > 0]} \end{cases} \quad (47)$$

and let $\pi^{(m)}$ be its stationary distribution. Now, it is enough to show that $\pi^{(m)}$ is uniquely determined by \mathbf{q} . First, observe that the process on $\bar{\mathbb{Z}}$ described in Equation (47) is reducible and can be decomposed into two irreducible

classes given by \mathbb{Z} and $\{\infty\}$, respectively. Therefore, if $\pi^{(m)}(Z = \infty) = 0$ or 1 , then it is unique. Indeed, if $\pi^{(m)}(Z = \infty) = 0$, then Z is birth-death process on \mathbb{Z} only, and hence it has a unique stationary distribution. Otherwise, if $\pi^{(m)}(Z = \infty) = 1$, then it is trivially unique. Now, we distinguish between two cases depending upon whether $q_m - q_{m+1} \geq \lambda$ or not.

Note that if $q_m - q_{m+1} \geq \lambda$, then $\pi^{(m)}(Z \geq k) = 1$ for all $k \geq 0$. On $\bar{\mathbb{Z}}$ this shows that $\pi^{(m)}(Z = \infty) = 1$. Furthermore, if $q_m - q_{m+1} < \lambda$, we will show that $\pi^{(m)}(Z = \infty) = 0$. On the contrary, assume $\pi^{(m)}(Z = \infty) = \varepsilon \in (0, 1]$. Also, let $\hat{\pi}^{(m)}$ be the unique stationary distribution of the birth-death process in Equation (47) on \mathbb{Z} . Therefore,

$$\pi_{\mathbf{q}}(\mathcal{R}_m) = \pi^{(m)}(Z > 0) = (1 - \varepsilon)\hat{\pi}^{(m)}(Z > 0) + \varepsilon = (1 - \varepsilon)\frac{q_m - q_{m+1}}{\lambda} + \varepsilon.$$

Substituting into the differential form of the fluid Equation (21) at the given time t , we obtain that

$$\frac{dq_m(t)}{dt} = \lambda \left[(1 - \varepsilon)\frac{q_m - q_{m+1}}{\lambda} + \varepsilon \right] - (q_m - q_{m+1}) = -\varepsilon(q_m - q_{m+1}) + \lambda\varepsilon > 0,$$

where the last inequality follows since we are considering the case when $q_m - q_{m+1} < \lambda$. Now, because $q_m(t) = 1$, this leads to a contradiction for any $\varepsilon > 0$, and hence it must be the case that $\pi^{(m)}(Z = \infty) = 0$.

Therefore, for all $\mathbf{q} \in \mathcal{S}$, $\pi_{\mathbf{q}}$ is uniquely determined by \mathbf{q} . Furthermore, we can identify the expression for $\pi_{\mathbf{q}}(\mathcal{R}_i)$ as

$$\pi_{\mathbf{q}}(\mathcal{R}_i) = \begin{cases} \min\{(q_m - q_{m+1})/\lambda, 1\} & \text{for } i = m, \\ 1 - \min\{(q_m - q_{m+1})/\lambda, 1\} & \text{for } i = m + 1, \\ 0 & \text{otherwise,} \end{cases} \quad (48)$$

and hence $\pi_{\mathbf{q}(s)}(\mathcal{R}_i) = p_{i-1}(\mathbf{q}(s))$ as claimed. \square

4.2. Equivalence on Fluid Scale

Having proved Theorem 5, it suffices to prove the universality property stated in the next proposition. This will complete the proof of Theorem 1.

Proposition 9. *If $d(N) \rightarrow \infty$ as $N \rightarrow \infty$, then the JSQ($d(N)$) scheme and the ordinary JSQ policy have the same fluid limit.*

The proof of the above proposition uses the S-coupling results from Section 3 and consists of three steps:

(i) First, we show that if $n(N)/N \rightarrow 0$ as $N \rightarrow \infty$, then the MJSQ($n(N)$) scheme has the same fluid limit as the ordinary JSQ policy.

(ii) Then, we apply Corollary 1 to prove that as long as $n(N)/N \rightarrow 0$, any scheme from the class CJSQ($n(N)$) has the same fluid limit as the ordinary JSQ policy.

(iii) Next, using Propositions 4 and 5, we establish that if $d(N) \rightarrow \infty$, then for some $n(N)$ with $n(N)/N \rightarrow 0$, the JSQ($d(N)$) scheme and the JSQ($n(N), d(N)$) scheme have the same fluid limit. The proposition then follows by observing that the JSQ($n(N), d(N)$) scheme belongs to the class CJSQ($n(N)$).

Proof of Proposition 9. First, to show Claim (i) above, define $\bar{N} = N - n(N)$ and $\bar{\lambda}(\bar{N}) = \lambda(N)$. Observe that the MJSQ($n(N)$) scheme with N servers can be thought of as the ordinary JSQ policy with \bar{N} servers and arrival rate $\bar{\lambda}(\bar{N})$. Also, because $n(N)/N \rightarrow 0$,

$$\frac{\bar{\lambda}(\bar{N})}{\bar{N}} = \frac{\lambda(N)}{N - n(N)} \rightarrow \lambda \quad \text{as } \bar{N} \rightarrow \infty.$$

Furthermore, observe that the fluid limit of the JSQ policy in Theorem 5 as given by Equation (21) is characterized by the parameter λ only, and hence the fluid limit of the MJSQ($n(N)$) scheme is the same as that of the ordinary JSQ policy.

Second, observe from the fluid limit of the JSQ policy that if $\lambda < 1$, then for any buffer capacity $b \geq 1$, and any starting state, the fluid-scaled cumulative overflow is negligible, i.e., for any $t \geq 0$, $L^N(t)/N \xrightarrow{\mathbb{P}} 0$. Because the above fact is induced by the fluid limit only, the same holds for the MJSQ($n(N)$) scheme. Therefore, using the lower and upper bounds in Corollary 1 and the tail bound in Proposition 3, we obtain Claim (ii) above.

Finally, choose $n(N) = N/\sqrt{d(N)}$ and consider the JSQ($n(N)$, $d(N)$) scheme. Because $d(N) \rightarrow \infty$, it is clear that $n(N)/N \rightarrow 0$ as $N \rightarrow \infty$. Also, if $\Delta^N(T)$ denotes the cumulative number of times that the JSQ($d(N)$) scheme and JSQ($n(N)$, $d(N)$) scheme differ in decision up to time T , then Proposition 5 yields

$$\mathbb{P}(\Delta^N(T) \geq \varepsilon N \mid A^N(T)) \leq \frac{A^N(T)}{\varepsilon N} \left(1 - \frac{n(N)}{N}\right)^{d(N)} = \frac{A^N(T)}{\varepsilon N} \left(1 - \frac{1}{\sqrt{d(N)}}\right)^{d(N)}.$$

Because $\{A^N(T)/N\}_{N \geq 1}$ is a tight sequence of random variables, we have

$$\frac{A^N(T)}{\varepsilon N} \left(1 - \frac{1}{\sqrt{d(N)}}\right)^{d(N)} \xrightarrow{\mathbb{P}} 0 \text{ as } N \rightarrow \infty,$$

and hence, $\Delta^N(T)/N \xrightarrow{\mathbb{P}} 0$. Therefore, applying the ℓ_1 distance bound stated in Proposition 4, we obtain Claim (iii). The proof is then completed by observing that the JSQ($n(N)$, $d(N)$) scheme belongs to the class CJSQ($n(N)$). \square

Proof of Theorem 2. For any $\varepsilon > 0$, define

$$T_\varepsilon^N := \inf \{t : Q_1^{d(N)}(t) > (\lambda + \varepsilon)N\}.$$

Now, the proof consists of two main steps. First, we show that if $d(N) \geq \ell(N)/(1 - \lambda - \varepsilon)$ for some $\varepsilon > 0$, then there exists an $\varepsilon' > 0$, such that if for some $T > 0$, $\mathbb{P}(T_\varepsilon^N > T) \rightarrow 1$ as $N \rightarrow \infty$, then the number of times that the JSQ($d(N)$) scheme and the ordinary JSQ policy differ in decision in $[0, T]$ is $o_p(N)$. This then implies that up to such a time T , it is enough to consider the fluid limit of the ordinary JSQ policy with batch arrivals. Second, we show that if the conditions stated in Theorem 2 hold, then for any finite time $T > 0$, $\mathbb{P}(T_\varepsilon^N > T) \rightarrow 1$ as $N \rightarrow \infty$. This will complete the proof.

To prove the first part, consider the JSQ($d(N)$) scheme in case of batch arrivals. Choose $\varepsilon' = \varepsilon/2$, and assume that $T > 0$ is such that $\mathbb{P}(T_{\varepsilon'}^N > T) \rightarrow 1$ as $N \rightarrow \infty$. Let I_i denote the number of idle servers among $d(N)$ randomly chosen servers for the i -th batch arrival, and define $W^N(t)$ to be the cumulative number of tasks that have not been assigned to some idle server, up to time t . If $A^N(t)$ denotes the number of batch arrivals that occurred up to time t , then

$$W^N(t) = \sum_{i=1}^{A^N(t)} [\ell(N) - I_i]^+.$$

We show that $W^N(t)/N \xrightarrow{\mathbb{P}} 0$ for all $t \leq T$ for $d(N) = \ell(N)/(1 - \lambda - \varepsilon)$. Observe that I_i follows a hypergeometric distribution with sample size $d(N)$ and population size N containing $N - Q_1^N(t) \geq (1 - \lambda - \varepsilon/2)N$ successes. Define J_i to be distributed as $d(N) - I_i$. Then,

$$[\ell(N) - I_i]^+ = k \iff J_i = d(N) - \ell(N) + k.$$

Therefore, for $c = 1 - \lambda - \varepsilon/2$ we have

$$\mathbb{E}([\ell(N) - I_i]^+) = \sum_{k \geq 1} k \mathbb{P}(J_i = (1 - c)d(N) + k) \leq d(N) \mathbb{P}(J_i \geq (1 - c)d(N)).$$

Now, from Hoeffding (1963) and Luh and Pippenger (2014), we know

$$\mathbb{P}(J_i \geq (1 - c)d(N)) \leq \exp(-d(N)H(\lambda, c)),$$

where

$$H(\lambda, c) = (1 - c) \log \left(\frac{1 - c}{\lambda}\right) + c \log \left(\frac{c}{1 - \lambda}\right) > 0,$$

because $c < 1 - \lambda$. Therefore,

$$\begin{aligned} \mathbb{P}(W^N(t) > \varepsilon N) &\leq \frac{\mathbb{E}(W^N(t))}{\varepsilon N} \\ &\leq \frac{d(N)}{\varepsilon N} \times \frac{\lambda(N)t}{\ell(N)} \times \exp(-d(N)H(\lambda, c)) \\ &= O(\exp(-d(N)H(\lambda, c))). \end{aligned} \tag{49}$$

This implies that whenever $\ell(N) \rightarrow \infty$, if $d(N) = \ell(N)/(1 - \lambda - \varepsilon/2)$, then $W^N(t)$ is $o_P(N)$ for all $t \leq T$. Now, the analysis of the ordinary JSQ policy with batch arrivals in Theorem 7 below, up to time T , shows that the process $\{\mathbf{q}^N(t)\}_{0 \leq t \leq T}$ converges to the deterministic limit $\{\mathbf{q}(t)\}_{0 \leq t \leq T'}$, described by Equation (3).

Therefore, it is enough to show that any $T > 0$ satisfies the required criterion. This can be seen by observing that for any $T \geq 0$ and any $\varepsilon' > 0$,

$$\begin{aligned} \mathbb{P}(T_{\varepsilon'}^N \leq T) &\leq \mathbb{P}(T_{\varepsilon'/2}^N < T) \\ &\leq \mathbb{P}\left(\sup_{t \in [0, T]} Q_1^{d(N)}(t) > (\lambda + \varepsilon'/2)N\right) \\ &\leq \mathbb{P}\left(\sup_{t \in [0, T]} Q_1^{\text{JSQ}}(t) > (\lambda + \varepsilon'/4)N\right) \times \mathbb{P}\left(\sup_{t \in [0, T]} |Q_1^{\text{JSQ}}(t) - Q_1^{d(N)}(t)| \leq \frac{N\varepsilon'}{4}\right) \\ &\quad + \mathbb{P}\left(\sup_{t \in [0, T]} |Q_1^{\text{JSQ}}(t) - Q_1^{d(N)}(t)| > \frac{N\varepsilon'}{4}\right) \rightarrow 0 \quad \text{as } N \rightarrow \infty. \end{aligned}$$

Therefore, the proof is complete. \square

Theorem 7 (Batch Arrivals JSQ). Consider the batch arrival scenario with growing batch size $\ell(N) \rightarrow \infty$ and $\lambda(N)/N \rightarrow \lambda < 1$ as $N \rightarrow \infty$. For the JSQ policy, if $q_1^{d(N)}(0) \xrightarrow{\mathbb{P}} q_1^\infty \leq \lambda$ and $q_i^{d(N)}(0) \xrightarrow{\mathbb{P}} 0$ for all $i \geq 2$, then the sequence of processes $\{\mathbf{q}^{d(N)}(t)\}_{t \geq 0}$ converges weakly to the limit $\{\mathbf{q}(t)\}_{t \geq 0}$ described as follows:

$$q_1(t) = \lambda + (q_1^\infty - \lambda)e^{-t}, \quad q_i(t) \equiv 0 \quad \text{for all } i = 2, \dots, b. \tag{50}$$

Proof. Fix any finite time $T \geq 0$. To analyze the JSQ policy with batch arrivals, observe that before time T , all the arriving tasks join idle servers. Therefore, assuming $Q_2^N(0) = 0$, for all $t \leq T$, the evolution for Q_1^N can be written as

$$Q_1^N(t) = Q_1^N(0) + \ell(N)A(t\lambda(N)/\ell(N)) - D\left(\int_0^t Q_1^N(s)ds\right), \tag{51}$$

where A and D are independent unit-rate Poisson processes. Using the random time change of unit-rate Poisson processes (Pang et al. 2007, lemma 3.2), and applying the arguments by Pang et al. (2007, lemma 3.4), the above process scaled by N then admits the martingale decomposition

$$q_1^N(t) = q_1^N(0) + \frac{M_1^N(t)}{N} + \lambda t - \frac{M_2^N(t)}{N} - \int_0^t q_1^N(s)ds, \tag{52}$$

where

$$\begin{aligned} M_1^N(t) &= \ell(N)A(t\lambda(N)/\ell(N)) - t\lambda(N), \\ M_2^N(t) &= D\left(\int_0^t Q_1^N(s)ds\right) - \int_0^t Q_1^N(s)ds, \end{aligned}$$

are square integrable martingales with respective quadratic variation processes given by

$$\begin{aligned} \langle M_1^N \rangle(t) &= t\lambda(N), \\ \langle M_2^N \rangle(t) &= \int_0^t Q_1^N(s)ds. \end{aligned}$$

Now, because for any $T \geq 0$, $\langle M_1^N \rangle(T)/N^2 \rightarrow 0$, and $\langle M_2^N \rangle(T)/N^2 \xrightarrow{\mathbb{P}} 0$ from the stochastic boundedness criterion for square integrable martingales (Pang et al. 2007, lemma 5.8), we get that both $\{M_1^N(t)/N\}_{t \geq 0} \xrightarrow{d} 0$ and $\{M_2^N(t)/N\}_{t \geq 0} \xrightarrow{d} 0$. Therefore, from the continuous mapping theorem and Equation (52), it follows that $\{q_1^N(t)\}_{t \geq 0}$ as $N \rightarrow \infty$ converges weakly to a deterministic limit described by the integral equation

$$q_1(t) = q_1^\infty + \lambda t - \int_0^t q_1(s)ds \tag{53}$$

having Equation (3) as the unique solution. This completes the proof of the fluid limit of the JSQ policy with batch arrivals. \square

4.3. Global Stability and Interchange of Limits

To prove the interchange of limits result stated in Proposition 1, we will establish the global stability of the fixed point, i.e., all fluid limit paths converge to the fixed point in Equation (2) as $t \rightarrow \infty$. This is formally stated in the following lemma.

Lemma 4. *Let $\mathbf{q}(t)$ be the fluid limit, i.e., the solution of the dynamical system described by the system of integral equations in Equation (1). For any $\mathbf{q}^\infty \in \mathcal{S}$, if $\mathbf{q}(0) = \mathbf{q}^\infty$, then $\mathbf{q}(t) \rightarrow \mathbf{q}^*$ as $t \rightarrow \infty$, where \mathbf{q}^* is defined as in Equation (2).*

In case of the JSQ(d) scheme with fixed d , the global stability is proved by constructing a Lyapunov function that measures the “distance” (in terms of a weighted L_1 -norm) between the trajectory and the fixed point and that strictly decreases everywhere except at the fixed point (see Mitzenmacher 1996, theorem 3.6). In case of the ordinary JSQ policy, however, we can exploit a more direct method to establish the global stability, as further detailed below.

Proof of Lemma 4. The proof follows in two steps: We will first establish that as $t \rightarrow \infty$, $q_1(t) \rightarrow \lambda < 1$ and then show that $q_2(t) \rightarrow 0$.

Observe that the rate of change of $q_1(t)$ is $\lambda p_0(\mathbf{q}(t)) - (q_1(t) - q_2(t))$. For any $\varepsilon \geq 0$ if $q_1(t) \leq \lambda - \varepsilon$, then $p_0(\mathbf{q}(t)) = 1$, so that the rate of change is $\lambda - (q_1(t) - q_2(t)) \geq \varepsilon$, i.e., positive and bounded away from zero when $\varepsilon > 0$. Also, $q_1(t)$ cannot decrease if $q_1(t) \leq \lambda$. This shows that for all $\varepsilon > 0$, there exists a time $t_0 = t_0(\varepsilon, \mathbf{q}^\infty)$ such that $q_1(t) \geq \lambda - \varepsilon$ for all $t \geq t_0$. Thus, $\liminf_{t \rightarrow \infty} q_1(t) \geq \lambda$.

On the other hand, we claim that $\limsup_{t \rightarrow \infty} q_1(t) \leq \lambda$. Suppose not, i.e., assume $\limsup_{t \rightarrow \infty} q_1(t) = \lambda + \varepsilon$ for some $\varepsilon > 0$. Because $q_1(t)$ is nondecreasing when $q_1(t) \leq \lambda$, there must exist a t_0 such that $q_1(t) \geq \lambda \forall t \geq t_0$. The high-level idea behind the claim is as follows. If $q_1(t)$ were to remain above λ by a nonvanishing margin, then the cumulative number of departures would exceed the cumulative number of arrivals by an infinite amount, which cannot occur because the initial number of tasks is bounded. More formally,

$$\begin{aligned} \sum_{i=1}^b q_i(t) &= \sum_{i=1}^b q_i(t_0) + \lambda \int_{t_0}^t \sum_{i=1}^b p_{i-1}(\mathbf{q}(s)) ds - \int_{t_0}^t q_1(s) ds \\ &\leq \sum_{i=1}^b q_i(t_0) - \int_{t_0}^t [q_1(s) - \lambda]^+ ds, \end{aligned}$$

and thus,

$$\int_{t_0}^t [q_1(s) - \lambda]^+ ds \leq \sum_{i=1}^b q_i(t) - \sum_{i=1}^b q_i(t_0) < \infty.$$

This provides a contradiction with $\limsup_{t \rightarrow \infty} q_1(t) = \lambda + \varepsilon$ because the rate of decrease of $q_1(t)$ is at most 1. Therefore, $q_1(t) \rightarrow \lambda$ as $t \rightarrow \infty$.

Consequently, for any $\mathbf{q}^\infty \in \mathcal{S}$ and $\varepsilon > 0$, if $\mathbf{q}(0) = \mathbf{q}^\infty$, then there exists a time $t_2 = t_2(\mathbf{q}^\infty, \varepsilon) < \infty$, such that $q_1(t) \leq \lambda + \varepsilon$ for all $t \geq t_2$. Thus, choosing $\varepsilon = (1 - \lambda)/2$ say, for all $t \geq t_2$, $q_1(t) < 1$, and thus $p_0(\mathbf{q}(t)) = 1$, i.e., $\sum_{i=2}^b p_{i-1}(\mathbf{q}(t)) = 0$. Define $q_{2+}(t) := \sum_{i=2}^b q_i(t)$. Observe that

$$\begin{aligned} q_{2+}(t) &= q_{2+}(t_2) + \lambda \int_{t_2}^t \sum_{i=2}^b p_{i-1}(\mathbf{q}(s)) ds - \int_{t_2}^t q_2(s) ds \\ &= q_{2+}(t_2) - \int_{t_2}^t q_2(s) ds \text{ for all } t \geq t_2, \end{aligned}$$

which implies $q_2(t) \leq q_{2+}(t_2)e^{-(t-t_2)}$. Thus, $q_2(t)$ and consequently, $q_{2+}(t)$ converges to 0 as $t \rightarrow \infty$. This completes the proof of global stability of the fixed point. \square

Proof of Proposition 1. The proof follows in two steps: (i) We first establish that the sequence of stationary measures $\{\pi^{d(N)}\}_{N \geq 1}$ is tight and then (ii) show the interchange of limits.

(i) Observe that if $b < \infty$, then the space $[0, 1]^b$ is compact, and hence Prohorov’s theorem implies that $\{\pi^{d(N)}\}_{N \geq 1}$ is tight. Now, assume $b = \infty$. For any two positive integers $d_1 \leq d_2$, note that at each arrival, the JSQ(d_2) scheme polls more servers than the JSQ(d_1) scheme. Thus, using the S-coupling and Proposition 2, we can conclude for every N ,

$$\sum_{i \geq m} Q_i^{d_2} \leq_{st} \sum_{i \geq m} Q_i^{d_1}, \text{ for all } m \geq 1.$$

In particular, putting $d_1 = 1$ and $d_2 = d(N)$,

$$\sum_{i \geq m} Q_i^{d(N)} \leq \sum_{i \geq m} Q_i^1, \quad \text{for all } m \geq 1. \tag{54}$$

Let \mathbf{X}^N and \mathbf{Y}^N denote random variables following the stationary distribution of two systems with N servers under the JSQ($d(N)$) and JSQ(1) schemes, respectively. We will verify the tightness criterion stated in Lemma 2. Note that because \mathbf{X}^N takes value in $\mathcal{S} \subset [0, 1]^\infty$, which is compact with respect to the product topology, Prohorov’s theorem implies that $\{\mathbf{X}^N\}_{N \geq 1}$ is tight with respect to the product topology. To verify the condition in Equation (39), note that the system under the JSQ(1) scheme is essentially a collection of N independent M/M/1 systems. Therefore, for each $k \geq 1$,

$$\overline{\lim}_{N \rightarrow \infty} \mathbb{P}\left(\sum_{i \geq k} X_i^N > \varepsilon\right) \leq \overline{\lim}_{N \rightarrow \infty} \mathbb{P}\left(\sum_{i \geq k} Y_i^N > \varepsilon\right) = (1 - \lambda) \sum_{i \geq k} \lambda^i.$$

Because $\lambda < 1$, taking the limit $k \rightarrow \infty$, the right side of the above inequality tends to zero, and hence, the condition in Equation (39) is verified.

(ii) Now, observe that because $\{\pi^{d(N)}\}_{N \geq 1}$ is tight, any subsequence has a convergent further subsequence. Let $\{\pi^{d(N_n)}\}_{n \geq 1}$ be any such convergent subsequence, with $\{N_n\}_{n \geq 1} \subseteq \mathbb{N}$, such that $\pi^{d(N_n)} \xrightarrow{d} \hat{\pi}$ as $n \rightarrow \infty$. We will show that $\hat{\pi}$ is unique and equals the measure π^* , as defined in the statement of Proposition 1. Notice that if $\mathbf{q}^{d(N_n)}(0) \sim \pi^{d(N_n)}$, then $\mathbf{q}^{d(N_n)}(t) \sim \pi^{d(N_n)}$ for all $t \geq 0$. Thus, $\hat{\pi}$ is an invariant distribution of the deterministic process $\{\mathbf{q}(t)\}_{t \geq 0}$. This in conjunction with the global stability in Lemma 4 implies that $\hat{\pi}$ must be the fixed point of the fluid limit. Thus, we have shown the convergence of the stationary measure. \square

5. Diffusion-Limit Proofs

In this section, we prove the diffusion-limit results for the JSQ($d(N)$) scheme stated in Theorem 3 and the almost necessity condition for diffusion-level optimality stated in Theorem 4. As noted in Subsection 2.4, the diffusion limit for the ordinary JSQ policy is obtained by Eschenfeldt and Gamarnik (2018, theorem 2) and characterized by Equation (4). Therefore, it suffices to prove the universality property stated in the next proposition.

Proposition 10. *If $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$ as $N \rightarrow \infty$, then the JSQ ($d(N)$) scheme and the ordinary JSQ policy have the same diffusion limit.*

The proof of the above proposition follows similar lines as that of Proposition 9, leveraging again the S-coupling results from Section 3, and involves three steps:

- (i) First, we show that if $n(N)/\sqrt{N} \rightarrow 0$ as $N \rightarrow \infty$, then the MJSQ($n(N)$) scheme has the same diffusion limit as the ordinary JSQ policy.
- (ii) Then, we use Corollary 1 to prove that as long as $n(N)/\sqrt{N} \rightarrow 0$, any scheme from the class CJSQ($n(N)$) has the same diffusion limit as the ordinary JSQ policy.
- (iii) Next, we establish using Propositions 4 and 5 that if $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$ as $N \rightarrow \infty$, then for some $n(N)$ with $n(N)/\sqrt{N} \rightarrow 0$, the JSQ($d(N)$) scheme and the JSQ($n(N), d(N)$) scheme have the same diffusion limit. The proposition then follows by observing that the JSQ($n(N), d(N)$) scheme belongs to the class CJSQ($n(N)$).

Proof of Proposition 10. To show Claim (i) above, define $\bar{N} = N - n(N)$ and $\bar{\lambda}(\bar{N}) = \lambda(N)$. As mentioned earlier, the MJSQ($n(N)$) scheme with N servers can be thought of as the ordinary JSQ policy with \bar{N} servers and arrival rate $\bar{\lambda}(\bar{N})$. Also, because $n(N)/\sqrt{N} \rightarrow 0$,

$$\frac{\bar{N} - \bar{\lambda}(\sqrt{\bar{N}})}{\bar{N}} = \frac{N - n(N) - \lambda(N)}{\sqrt{N} - n(N)} \rightarrow \beta > 0 \quad \text{as } \bar{N} \rightarrow \infty.$$

Furthermore, observe that the diffusion limit of the JSQ policy of Eschenfeldt and Gamarnik (2018, theorem 2) as given in Equation (4) is characterized by the parameter $\beta > 0$, and hence the diffusion limit of the MJSQ($n(N)$) scheme is the same as that of the ordinary JSQ policy.

Observe from the diffusion limit of the JSQ policy that if $\beta > 0$, then for any buffer capacity $b \geq 2$, and suitable initial state as described in Theorem 3, the cumulative overflow is negligible, i.e., for any $t \geq 0$, $L^N(t) \xrightarrow{\mathbb{P}} 0$. Indeed, observe that if $b \geq 2$, and $\{\bar{Q}_2^N(0)\}_{N \geq 1}$ is a tight sequence, then the sequence of processes $\{\bar{Q}_2^N(t)\}_{t \geq 0}$ is stochastically bounded. Therefore, on any finite time interval, there will be only $O_P(\sqrt{N})$ servers with queue length larger than

one; whereas, for an overflow event to occur, all the N servers must have at least two pending tasks. Therefore, for any $t \geq 0$,

$$\begin{aligned} \limsup_{N \rightarrow \infty} \mathbb{P}(L^N(t) > 0) &\leq \limsup_{N \rightarrow \infty} \mathbb{P}\left(\sup_{s \in [0, t]} Q_2^N(s) = N\right) \\ &\leq \limsup_{N \rightarrow \infty} \mathbb{P}\left(\sup_{s \in [0, t]} \bar{Q}_2^N(s) = \sqrt{N}\right) = 0. \end{aligned}$$

Because the above fact is implied by the diffusion limit only, the same holds for the MJSQ($n(N)$) scheme. Therefore, using the lower and upper bounds in Corollary 1 we arrive at Claim (ii).

Finally, choose

$$n(N) = \frac{N \log N}{d(N)},$$

and consider the JSQ($n(N), d(N)$) scheme. Because $d(N)/(\sqrt{N} \log N) \rightarrow \infty$, it is clear that $n(N)/\sqrt{N} \rightarrow 0$ as $N \rightarrow \infty$. Again, if $\Delta^N(T)$ denotes the cumulative number of times that the JSQ($d(N)$) scheme and JSQ($n(N), d(N)$) scheme differ in decision up to time T , then Proposition 5 yields

$$\begin{aligned} \mathbb{P}\left(\Delta^N(T) \geq \varepsilon \sqrt{N} \mid A^N(T)\right) &\leq \frac{A^N(T)}{\varepsilon \sqrt{N}} \left(1 - \frac{n(N)}{N}\right)^{d(N)} \\ &\leq \frac{A^N(T)}{\varepsilon \sqrt{N}} \left(1 - \frac{\log(N)}{d(N)}\right)^{d(N)} \\ &\leq \frac{A^N(T)}{\varepsilon N} \sqrt{N} \left(1 - \frac{\log(N)}{d(N)}\right)^{d(N)}. \end{aligned} \tag{55}$$

Because $\{A^N(T)/N\}_{N \geq 1}$ is a tight sequence of random variables and

$$\begin{aligned} \sqrt{N} \left(1 - \frac{\log(N)}{d(N)}\right)^{d(N)} &\rightarrow 0, \quad \text{as } N \rightarrow \infty, \\ \iff \frac{1}{2} \log(N) + d(N) \log \left(1 - \frac{\log(N)}{d(N)}\right) &\rightarrow -\infty, \quad \text{as } N \rightarrow \infty, \\ \iff \frac{1}{2} \log N - \frac{\log(N)}{d(N)} \times d(N) &\rightarrow -\infty, \quad \text{as } N \rightarrow \infty, \end{aligned}$$

from Equation (55), $\Delta^N(T)/N \xrightarrow{\mathbb{P}} 0$. Therefore, by invoking Proposition 4, we obtain Claim (iii). The proof is then completed by observing that the JSQ($n(N), d(N)$) scheme belongs to the class CJSQ($n(N)$). \square

We next prove that the growth condition $d(N)/(\sqrt{N} \log N) \rightarrow \infty$ is nearly necessary: For any $d(N)$ such that $d(N)/(\sqrt{N} \log N) \rightarrow 0$ as $N \rightarrow \infty$, the diffusion limit of the JSQ($d(N)$) scheme differs from that of the ordinary JSQ policy. Note that it is enough to consider the truncated system where any arrival to a server with at least two tasks is discarded, because the truncated system and the original system have the same diffusion limit (Mukherjee et al. 2016b).

Now, consider the JSQ($d(N)$) scheme for some $d(N)$ such that $d(N)/(\sqrt{N} \log N) \rightarrow 0$ as $N \rightarrow \infty$, and assume on the contrary, the hypothesis that the process $\{(Q_1^{d(N)}(t) - N)/\sqrt{N}, Q_2^{d(N)}(t)/\sqrt{N}\}_{t \geq 0}$ converges to the diffusion limit corresponding that of the JSQ policy. From a high level, the idea is to show that if the processes $(N - Q_1^{d(N)}(\cdot))$ and $Q_2^{d(N)}(\cdot)$ are $O_p(\sqrt{N})$, then in any finite time interval the number of tasks assigned to a server with queue length at least one, by the JSQ($d(N)$) scheme with $d(N)/(\sqrt{N} \log N) \rightarrow 0$ does not scale with \sqrt{N} , which then immediately proves that the diffusion limit cannot coincide with that of the ordinary JSQ policy.

To formalize the above idea, we first define an artificial scheme below, which will serve as an asymptotic lower bound to the number of servers with queue length two in a system following the JSQ($d(N)$) scheme under the hypothesis that the diffusion limit of the JSQ($d(N)$) coincides with that of the ordinary JSQ policy. For any nonnegative sequence $c(N)$, define a scheme $\Pi(c(N))$, which

(i) At each external arrival, assigns the task to a server having queue length one with probability $(1 - c(N)/N)^{d(N)}$ and otherwise discards it (ties can be broken randomly).

(ii) If a departure occurs from a server with queue length one, then it immediately makes the server busy with a dummy arrival, i.e., essentially $\Pi(c(N))$ prohibits any server to remain idle.

We use a coupling argument to show the following:

Lemma 5. *For any nonnegative sequence $c(N)$ with $c(N)/\sqrt{N} \rightarrow \infty$ as $N \rightarrow \infty$, there exists a common probability space such that for any $T > 0$,*

$$\mathbb{P} \left(\sup_{t \in [0, T]} \left\{ Q_2^{d(N)}(t) - Q_2^{\Pi(c(N))}(t) \right\} \geq 0 \right) \rightarrow 1 \quad \text{as } N \rightarrow \infty,$$

provided $Q_2^{d(N)}(0) \geq Q_2^{\Pi(c(N))}(0)$ for all sufficiently large N and the sequence of processes $\{(N - Q_1^{d(N)}(t))/\sqrt{N}\}_{t \geq 0}$ and $\{Q_2^{d(N)}(t)/\sqrt{N}\}_{t \geq 0}$ are stochastically bounded.

In order to prove Lemma 5, we first S-couple the two systems under schemes $\Pi(c(N))$ and JSQ($d(N)$), respectively. Now, at each external arrival, to assign the task in the two systems in a coupled way, draw a single uniform $[0, 1]$ random variable U , independent of any other processes.

• Under the JSQ($d(N)$) scheme, if $u < (Q_1^{d(N)}/N)^{d(N)} - (Q_2^{d(N)}/N)^{d(N)}$, then assign the task to a server with queue length one, if

$$(Q_1^{d(N)}/N)^{d(N)} - (Q_2^{d(N)}/N)^{d(N)} < U < 1 - (Q_2^{d(N)}/N)^{d(N)}, \quad (56)$$

then assign the task to an idle server, and otherwise discard it. This preserves the statistical law of the JSQ($d(N)$) scheme with a buffer size $b = 2$. Indeed note that according to the above rule the probability that an incoming task will be assigned to some server with queue length zero, one, and two, are, respectively given by $(1 - (Q_1^{d(N)}/N)^{d(N)})$, $(Q_1^{d(N)}/N)^{d(N)} - (Q_2^{d(N)}/N)^{d(N)}$, and $(Q_2^{d(N)}/N)^{d(N)}$.

• Under the scheme $\Pi(c(N))$, if $U < (1 - c(N)/N)^{d(N)}$, then assign the incoming task to a server with queue length one, otherwise discard it. Clearly, the statistical law of the $\Pi(c(N))$ scheme is preserved by this rule.

Proof of Lemma 5. Fix any $T \geq 0$. Now, the proof follows in two steps:

(i) First, assume that at each external arrival up to time T , whenever an incoming task joins a server with queue length one, under the $\Pi(c(N))$ scheme, then so does the incoming task under JSQ($d(N)$) scheme. In that case, because the two systems are S-coupled, by forward induction on event times, it can be seen that $Q_2^{d(N)}(t) \geq Q_2^{\Pi(c(N))}(t)$ for all $0 \leq t \leq T$, provided $Q_2^{d(N)}(0) \geq Q_2^{\Pi(c(N))}(0)$.

(ii) Now, for any $T \geq 0$, according to the hypothesis, both $\sup_{t \in [0, T]} Q_2^{d(N)}(t)$ and $\sup_{t \in [0, T]} \{N - Q_1^{d(N)}(t)\}$ are $O_P(\sqrt{N})$. Also, because $c(N)/\sqrt{N} \rightarrow \infty$, it is straightforward to check that

$$\liminf_{N \rightarrow \infty} \mathbb{P} \left(\sup_{t \in [0, T]} \left(\frac{Q_1^{d(N)}(t)}{N} \right)^{d(N)} - \left(\frac{Q_2^{d(N)}(t)}{N} \right)^{d(N)} \geq \left(1 - \frac{c(N)}{N} \right)^{d(N)} \right) = 1. \quad (57)$$

Note that the probabilities that an incoming task joins a server with queue length one are given by $(Q_1^{d(N)}(t)/N)^{d(N)} - (Q_2^{d(N)}(t)/N)^{d(N)}$ and $(1 - c(N)/N)^{d(N)}$ for the JSQ($d(N)$) and the $\Pi(c(N))$ scheme, respectively. Informally speaking, because of the above coupling, Equation (57) then implies that with high probability, on any finite time interval, whenever an external incoming task joins a server with queue length one under the $\Pi(c(N))$ scheme, then so does the incoming task under the JSQ($d(N)$) scheme. Therefore, from Part (i) above, we obtain

$$\begin{aligned} & \mathbb{P} \left(\sup_{t \in [0, T]} \left\{ Q_2^{d(N)}(t) - Q_2^{\Pi(c(N))}(t) \right\} \geq 0 \right) \\ & \geq \mathbb{P} \left(\sup_{t \in [0, T]} \left(\frac{Q_1^{d(N)}(t)}{N} \right)^{d(N)} - \left(\frac{Q_2^{d(N)}(t)}{N} \right)^{d(N)} \geq \left(1 - \frac{c(N)}{N} \right)^{d(N)} \right) \\ & \rightarrow 1 \quad \text{as } N \rightarrow \infty. \end{aligned}$$

This shows that if for any $T \geq 0$, if $\sup_{t \in [0, T]} Q_2^{d(N)}(t)$ and $\sup_{t \in [0, T]} \{N - Q_1^{d(N)}(t)\}$ are $O_P(\sqrt{N})$, then with probability tending to one as $N \rightarrow \infty$, up to time t , the process $\{Q_2^{\Pi(c(N))}(t)\}_{0 \leq t \leq T}$ is, indeed, a lower bound for the process $\{Q_2^{d(N)}(t)\}_{0 \leq t \leq T}$, and hence by our hypothesis, the proof is complete. \square

Proof of Theorem 4. Fix any positive sequence $d(N)$ such that $d(N)/(\sqrt{N} \log N) \rightarrow 0$ as $N \rightarrow \infty$. Assume that for the JSQ($d(N)$) scheme, the process

$$\{(Q_1^{d(N)}(t) - N)/\sqrt{N}, Q_2^{d(N)}(t)/\sqrt{N}\}_{t \geq 0}$$

converges to the appropriate diffusion limit corresponding to that of the ordinary JSQ policy. We will show that under this hypothesis, the process $\{Q_2^{d(N)}(t)/\sqrt{N}\}_{t \geq 0}$ is not stochastically bounded, which will then lead to a contradiction.

In order to show this, we will choose an appropriate $c(N)$ such that $c(N)/\sqrt{N} \rightarrow \infty$ as $N \rightarrow \infty$, and the process $\{Q_2^{\Pi(c(N))}(t)/\sqrt{N}\}_{t \geq 0}$ is not stochastically bounded. The conclusion then follows by the application of Lemma 5.

Observe that the martingale decomposition of the scaled $Q_2^{\Pi(c(N))}(\cdot)$ process can be written as

$$\bar{Q}_2^{\Pi(c(N))}(t) = \bar{Q}_2^{\Pi(c(N))}(0) + \frac{M^N(t)}{\sqrt{N}} + \frac{\lambda(N)t}{\sqrt{N}} \left(1 - \frac{c(N)}{N}\right)^{d(N)} - \int_0^t \bar{Q}_2^{\Pi(c(N))}(s) ds, \quad (58)$$

where $\bar{Q}_2^{\Pi(c(N))}(t) = Q_2^{\Pi(c(N))}(t)/\sqrt{N}$. Now, write $c(N) = g(N)\sqrt{N}$, for some $g(N) \rightarrow \infty$ (to be chosen later), and $d(N) = \sqrt{N} \log(N)/\omega(N)$, where $\omega(N) = \sqrt{N} \log(N)/d(N) \rightarrow \infty$ as $N \rightarrow \infty$. Therefore, we write Equation (58) as

$$\bar{Q}_2^{\Pi(c(N))}(t) = \bar{Q}_2^{\Pi(c(N))}(0) + \frac{M^N(t)}{\sqrt{N}} + \frac{\lambda(N)t}{\sqrt{N}} \left(1 - \frac{g(N)}{\sqrt{N}}\right)^{\frac{\sqrt{N} \log N}{\omega(N)}} - \int_0^t \bar{Q}_2^{\Pi(c(N))}(s) ds. \quad (59)$$

Observe that for any $t \geq 0$,

$$\begin{aligned} & \lim_{N \rightarrow \infty} \frac{\lambda(N)t}{\sqrt{N}} \left(1 - \frac{g(N)}{\sqrt{N}}\right)^{\frac{\sqrt{N} \log N}{\omega(N)}} \\ &= t \lim_{N \rightarrow \infty} \exp \left[\log(\sqrt{N} - \beta) + \frac{\sqrt{N} \log N}{\omega(N)} \log \left(1 - \frac{g(N)}{\sqrt{N}}\right) \right] \\ &= t \lim_{N \rightarrow \infty} \exp \left[\log(\sqrt{N} - \beta) - \frac{g(N) \log N}{\omega(N)} - o\left(\frac{g(N) \log N}{\omega(N)}\right) \right]. \end{aligned}$$

Choosing $g(N)$ such that $g(N)/\omega(N) \rightarrow 0$ implies

$$\frac{\lambda(N)t}{\sqrt{N}} \left(1 - \frac{g(N)}{\sqrt{N}}\right)^{\frac{\sqrt{N} \log N}{\omega(N)}} \rightarrow \infty, \quad \text{as } N \rightarrow \infty.$$

Note that for any $\omega(N)$, this choice of $g(N)$ is feasible (choose $g(N) = \sqrt{\omega(N)}$, say). Furthermore, the process $\{M^N(t)/\sqrt{N}\}_{t \geq 0}$ in Equation (58) is stochastically bounded due to the martingale FCLT (Ethier and Kurtz 2009, theorem 7.1) and our hypothesis. Now, we can conclude that for the above choices of $g(N)$ and $\omega(N)$, the process $\{\bar{Q}_2^{\Pi(c(N))}(t)\}_{t \geq 0}$, and hence the process $\{Q_2^{d(N)}(t)\}_{t \geq 0}$ (because of Lemma 5) is not stochastically bounded. Therefore, the limit does not coincide with the limit of the scaled Q_2^{JSQ} -process. \square

6. Conclusion

In the present paper, we have established universality properties for power-of- d load balancing schemes in many-server systems. Specifically, we considered a system of N parallel exponential servers and a single dispatcher that assigns arriving tasks to the server with the shortest queue among $d(N)$ randomly selected servers. We developed a novel stochastic coupling construction to bound the difference in the queue length processes between the JSQ policy ($d = N$) and a scheme with an arbitrary value of d . As it turns out, a direct comparison between the JSQ policy and a JSQ(d) scheme is a significant challenge. Hence, we adopted a two-stage approach based on a novel class of schemes that always assign the incoming task to one of the servers with the $n(N) + 1$ smallest queue lengths. Just like the JSQ($d(N)$) scheme, these schemes may be thought of as “sloppy” versions of the JSQ policy. Indeed, the JSQ($d(N)$) scheme is guaranteed to identify the server with the minimum queue length but only among a randomly sampled subset of $d(N)$ servers. In contrast, the schemes in the above class only guarantee that one of the $n(N) + 1$ servers with the smallest queue lengths is selected, but across the entire system of N servers. We showed that the system occupancy processes for an intermediate blend of these schemes are simultaneously close on a $g(N)$ scale ($g(N) = N$ or $g(N) = \sqrt{N}$) to both the JSQ policy and the JSQ($d(N)$) scheme for suitably chosen values of $d(N)$ and $n(N)$ as function of $g(N)$. Given the latter asymptotic universality, it then sufficed to establish the fluid and

diffusion limits for the ordinary JSQ policy. Thus, deriving the fluid limit of the ordinary JSQ policy and using the above coupling argument, we establish the fluid limit of the JSQ($d(N)$) scheme in a regime with $d(N) \rightarrow \infty$ as $N \rightarrow \infty$, along with the corresponding fixed point. The fluid limit turns out not to depend on the exact growth rate of $d(N)$, and, in particular, coincides with that for the ordinary JSQ policy. We further leveraged the coupling to prove that the diffusion limit in the Halfin-Whitt regime with $d(N)/(\sqrt{N} \log(N)) \rightarrow \infty$ as $N \rightarrow \infty$ corresponds to that for the JSQ policy. These results indicate that the optimality of the JSQ policy can be preserved at the fluid level and diffusion level while reducing the overhead by nearly a factor $O(N)$ and $O(\sqrt{N}/\log(N))$, respectively. In future work we plan to extend the results to heterogeneous servers and nonexponential service requirement distributions. We also intend to pursue extensions to network scenarios and server-task compatibility constraints.

Acknowledgments

Debankur Mukherjee's current affiliation is Brown University, Providence, Rhode Island.

References

- Banerjee S, Mukherjee D (2018) Join-the-shortest queue diffusion limit in Halfin-Whitt regime: Tail asymptotics and scaling of extrema. *Ann. Appl. Probab.* Forthcoming. arXiv preprint available at <http://arxiv.org/abs/1803.03306>.
- Bramson M, Lu Y, Prabhakar B (2012) Asymptotic independence of queues under randomized load balancing. *Queueing Systems* 71(3):247–292.
- Braverman A (2018) Steady-state analysis of the join the shortest queue model in the Halfin-Whitt regime. Working paper, Northwestern University, Chicago. <https://arxiv.org/pdf/1801.05121.pdf>.
- Brightwell G, Fairthorne M, Luczak MJ (2018) The supermarket model with bounded queue lengths in equilibrium. *J. Statist. Phys.* 173(3–4): 1149–1194.
- Ephremides A, Varaiya P, Walrand J (1980) A simple dynamic routing problem. *IEEE Trans. Automatic Control* 25(4):690–693.
- Eschenfeldt P, Gamarnik D (2016) Supermarket queueing system in the heavy traffic regime. Short queue dynamics. Working paper, Harvard Medical School, Boston. <http://arxiv.org/abs/1610.03522>.
- Eschenfeldt P, Gamarnik D (2018) Join the shortest queue with many servers. The heavy traffic asymptotics. *Math. Oper. Res.* 43(3):867–886.
- Ethier SN, Kurtz TG (2009) *Markov Processes: Characterization and Convergence* (John Wiley & Sons, Hoboken, NJ).
- Gamarnik D, Tsitsiklis J, Zubeldia M (2016) Delay, memory and messaging tradeoffs in distributed service systems. *Proc. SIGMETRICS '16* (ACM, New York), 1–12.
- Graham C (2005) Functional central limit theorems for a large network in which customers join the shortest of several queues. *Probab. Theory Related Fields* 131(1):97–120.
- Gupta V, Harchol-Balter M, Sigman K, Whitt W (2007) Analysis of join-the-shortest-queue routing for web server farms. *Performance Evaluation* 64(9–12):1062–1081.
- Halfin S, Whitt W (1981) Heavy-traffic limits for queues with many exponential servers. *Oper. Res.* 29(3):567–588.
- Hoeffding W (1963) Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* 58(301):13–30.
- Hunt P, Kurtz T (1994) Large loss networks. *Stochastic Processes Their Appl.* 53(2):363–378.
- Jonckheere M (2006) Insensitive versus efficient dynamic load balancing in networks without blocking. *Queueing Systems* 54(3):193–202.
- Liptser R, Shiryaev A (1989) *Theory of Martingales* (Springer, New York).
- Luczak MJ, McDiarmid C (2006) On the maximum queue length in the supermarket model. *Ann. Probab.* 34(2):493–527.
- Luczak MJ, Norris J (2005) Strong approximation for the supermarket model. *Ann. Appl. Probab.* 15(3):2038–2061.
- Luh K, Pippenger N (2014) Large-deviation bounds for sampling without replacement. *Amer. Math. Monthly* 121(5):449–454.
- Mitzenmacher M (1996) The power of two choices in randomized load balancing. PhD thesis, University of California, Berkeley.
- Mitzenmacher M (2001) The power of two choices in randomized load balancing. *IEEE Trans. Parallel Distributed Systems* 12(10):1094–1104.
- Mukherjee D, Borst SC, Van Leeuwaarden JSH, Whiting PA (2016a) Asymptotic optimality of power-of- d load balancing in large-scale systems. Working paper, Brown University, Providence, RI. <https://arxiv.org/abs/1612.00722>.
- Mukherjee D, Borst SC, Van Leeuwaarden JSH, Whiting PA (2016b) Universality of load balancing schemes on the diffusion scale. *J. Appl. Probab.* 53(4):1111–1124.
- Pang G, Talreja R, Whitt W (2007) Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probab. Surveys* 4:193–267.
- Sparaggis PD, Towsley D, Cassandras CG (1994) Sample path criteria for weak majorization. *Adv. Appl. Probab.* 26(1):155–171.
- Towsley D (1995) Application of majorization to control problems in queueing systems. Chrétienne P, Coffman EG, Lenstra JK, Liu Z, eds. *Scheduling Theory and Its Applications* (John Wiley & Sons, Hoboken, NJ), 295–312.
- Towsley D, Sparaggis P, Cassandras C (1992) Optimal routing and buffer allocation for a class of finite capacity queueing systems. *IEEE Trans. Automatic Control* 37(9):1446–1451.
- Van der Boor M, Borst SC, Van Leeuwaarden JSH, Mukherjee D (2018) Scalable load balancing in networked systems: Universality properties and stochastic coupling methods. Sirakov B, de Souza PN, Viana M, eds. *Proc. ICM '18, Rio de Janeiro, Brazil* (World Scientific).
- Van Leeuwaarden JSH, Knessl C (2011) Transient behavior of the Halfin-Whitt diffusion. *Stochastic Processes Their Appl.* 121(7):1524–1545.
- Van Leeuwaarden JSH, Knessl C (2012) Spectral gap of the Erlang A model in the Halfin-Whitt regime. *Stochastic Systems* 2(1):149–207.
- Vvedenskaya ND, Dobrushin RL, Karpelevich FI (1996) Queueing system with selection of the shortest of two queues: An asymptotic approach. *Problemy Peredachi Informatsii* 32(1):20–34.
- Weber RR (1978) On the optimal assignment of customers to parallel servers. *J. Appl. Probab.* 15(2):406–413.
- Whitt W (1986) Deciding which queue to join: Some counterexamples. *Oper. Res.* 34(1):55–62.
- Whitt W (2002) *Stochastic-Process Limits*. Springer Series in Operations Research and Financial Engineering (Springer-Verlag, New York).
- Winston W (1977) Optimality of the shortest line discipline. *J. Appl. Probab.* 14(1):181–189.
- Ying L, Srikant R, Kang X (2015) The power of slightly more than one sample in randomized load balancing. *Proc. 2015 IEEE Conf. Comput. Comm., (INFOCOM)* (Institute of Electrical and Electronics Engineers Inc., Piscataway, NJ), 1131–1139.