



The Power of Two Choices with Load Comparison Errors

Sanidhay Bhambay
Department of Computer Science,
University of Warwick
UK

Arpan Mukhopadhyay
Department of Computer Science,
University of Warwick
UK

Thirupathaiah Vasantam
Department of Computer Science,
Durham University
UK

ABSTRACT

We consider a system with n unit-rate servers where jobs arrive according a Poisson process with rate $n\lambda$ ($\lambda < 1$). In the standard *Power-of-two* or Po2 scheme, for each incoming job, a job dispatcher samples two servers uniformly at random and sends the incoming job to the least loaded of the two sampled servers. However, in practice, the load information may not be accurate at the job dispatcher. In this paper, we analyze the effects of erroneous load comparisons on the performance of the Po2 scheme. Specifically, we consider *load-dependent* and *load-independent* errors. In the load-dependent error model, an incoming job is sent to the server with the larger queue length among the two sampled servers with an error probability ϵ if the difference in the queue lengths of the two sampled servers is less than or equal to a constant g ; no error is made if the queue-length difference is higher than g . For this type of errors, we show that, in the large system limit, the benefits of the Po2 scheme are retained for all values of g and ϵ as long as the system is heavily loaded, i.e., λ is close to 1. In the load-independent error model, the incoming job is sent to the sampled server with the *maximum load* with an error probability of ϵ independent of the loads of the sampled servers. For this model, we show that the performance benefits of the Po2 scheme are retained only if $\epsilon \leq 1/2$; for $\epsilon > 1/2$ we show that the stability region of the system reduces and the system performs poorly in comparison to the *random scheme*. To prove our stability results, we develop a generic approach to bound the drifts of Lyapunov functions for any state-dependent load balancing scheme. Furthermore, the mean-field analysis in our paper uses a new approach to characterise fixed points which do not admit a recursion.

CCS CONCEPTS

• Networks → Network performance analysis;

KEYWORDS

Mean-field analysis, Lyapunov drift analysis, Load balancing, Power-of-Two

ACM Reference Format:

Sanidhay Bhambay, Arpan Mukhopadhyay, and Thirupathaiah Vasantam. 2023. The Power of Two Choices with Load Comparison Errors. In *International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing (MobiHoc '23)*, October 23–26,

2023, Washington, DC, USA. ACM, New York, NY, USA, Article 111, 10 pages. <https://doi.org/10.1145/3565287.3610259>

1 INTRODUCTION

Modern data centres consist of large numbers of parallel servers. Balancing the load across these servers is crucial to ensure better resource utilization and satisfactory quality of service. To distribute the load across the servers uniformly, each incoming job needs to be assigned to an appropriate server in the system. Assigning an incoming job to one of the servers is done by a job dispatcher using a load balancing scheme. For homogeneous systems, it is well known that the *Join-the-Shortest-Queue* (JSQ) load balancing scheme [1, 2], where an incoming job is assigned to the server having smallest number of ongoing jobs, is optimal in terms of minimizing the average response time of jobs. For heterogeneous systems, variants of the JSQ scheme are known to be asymptotically optimal [3, 4]. However, implementing JSQ-type schemes is often difficult in practice as dispatchers only have local views of the system and therefore can access the load information of only a subset of servers [5, 6]. A popular alternative to the JSQ scheme, called the *Join-the-Idle-Queue* scheme [7, 8], requires only the knowledge of the idle servers. However, even this scheme suffers from similar implementation issues as the dispatcher needs to store the idle tokens of a large number of servers.

Due to the challenges mentioned above, randomised schemes, such as the celebrated *Power-of-d choices* or the *Pod* scheme [9, 10], are widely used in practice [11, 12]. In the Pod scheme, an arrival is sent to the server with the shortest queue length among a set of d servers sampled uniformly at random. For $d = 1$, the Pod scheme reduces to the *random scheme*. For $d = 2$, it is well-known that an exponential improvement in the average response time is achieved in comparison to the random scheme. However, this result requires that the server with the least queue length among the two sampled servers is always correctly identified. In a real system, the queue lengths stored at a dispatcher may get outdated due to infrequent updates from the servers [13, 14]. This could result in misidentifying the server currently having the smaller queue length. Another scenario where an error can occur, is when an adversary tries to carry out an attack by misreporting the queue-lengths sent from the servers to the dispatcher [15]. The attacker can manipulate the queue lengths in a way that the dispatcher assigns the job to the server with the maximum load among the two sampled servers. Such erroneous assignments can significantly increase average response time of jobs and may even cause the system to become unstable. The importance of studying the effect of inaccurate load comparisons on load balancing was highlighted as early as 2001 in a survey paper by Mitzenmacher, Richa and Sitaraman [16]. However, except from the static setting [17, 18], where there is a finite pool of jobs, this problem has not been studied.



This work is licensed under a Creative Commons Attribution International 4.0 License.

MobiHoc '23, October 23–26, 2023, Washington, DC, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9926-5/23/10.

<https://doi.org/10.1145/3565287.3610259>

Our Contributions: This motivates us to consider the effects that comparison errors can have on the performance of the Po2 scheme in the dynamic setting. We consider two types of comparison errors, referred to as the *load-dependent* and *load-independent* errors. In the load-dependent error model, an “error” is made with probability $\epsilon \in [0, 1]$ if the difference in the queue lengths of the sampled servers is sufficiently small (less or equal to a constant $g \geq 0$); if this is not the case (i.e., if the queue-length difference is higher than g), then no error occurs. An error, in this context, refers to the event where the job is sent to the server having the larger queue length among the two sampled servers. Hence, in this model, the dispatcher makes an error only when the current queue lengths of the sampled servers are close to each other; this is natural to expect when errors occur primarily due to outdated queue lengths at the dispatcher as servers having close queue lengths are likely to be more affected by this type of errors. To model errors due to adversarial attacks, we consider the load-independent error model in which an error occurs with probability ϵ independent of the current loads of the sampled servers. Clearly, this model of error can have a more drastic impact on the system’s performance than the load-dependent error model. Our goal is to characterise the performance of the Po2 scheme under these two error models for a system where there are n unit-rate servers and jobs with exponentially distributed sizes arrive according to a Poisson process with rate $n\lambda$ ($\lambda < 1$).

It is natural to expect that as g and ϵ increases (i.e., as the error rate becomes higher), the performance of the Po2 scheme under the load-dependent error model would deteriorate, eventually resulting in a performance poorer than the random scheme. While this is true for light traffic (small values of λ), we show that, in the heavy traffic limit ($\lambda \rightarrow 1$) and large system sizes, the performance of the system remains exponentially better than that under the random scheme for all values of g and ϵ . This implies that the benefits of sampling one additional server in the Po2 scheme outweighs the negative impact of comparison errors when the system operates at its maximum capacity. This result can be explained through the dynamics of the system in the mean-field regime. Specifically, we show the fixed point of the mean-field has a super-exponential decay of tail probabilities for all values of g and ϵ . While this decay rate is slower than that under the standard Po2 scheme, it is still super-exponential and therefore its benefits in comparison to the exponential decay rate under the random scheme become more prominent in the heavy traffic regime.

For the load-independent error model, we show that the benefits of the Po2 scheme are retained only if the error probability ϵ satisfies $\epsilon \leq 1/2$. For $\epsilon > 1/2$, we show that system becomes unstable for arrival rates larger than $1/2\epsilon$ and the performance becomes worse than that under the random scheme. This can be intuitively explained by the fact that for $\epsilon > 1/2$ the Po2 scheme chooses the server with the larger queue length more often than the server with the smaller queue length.

From a technical point of view, we make a number of important contributions. First, we derive the stability region of the system for both error models and establish uniform (in the system size n) bounds on the stationary expected queue length per server. These bounds are essential to establish tightness of the stationary measures and interchange of limits in the mean-field regime. The existing results of [19] on JSQ-type load balancing schemes do not apply

to our schemes since a job is not always sent to the server with the minimum queue length among the sampled servers. Although the fluid limit results of [20] can be applied to derive stability conditions, they do not yield the uniform bounds essential to establish tightness of the stationary distributions. To obtain both stability conditions and the uniform bounds, we use drifts of suitable Lyapunov functions. However, bounding the drifts of these Lyapunov functions is difficult for our schemes as the schemes compare only a subset of queues at each arrival and do not always choose the shortest queue as the final destination. We develop a generic approach through which the drift can be bounded for any scheme where queue lengths of multiple servers are compared to dispatch each job.

The second important technical contribution is the mean-field analysis of the Po2 scheme under the load-dependent error model. This analysis differs significantly from conventional analysis in that the fixed point of the mean-field in this case does not satisfy any recursion. For such a system, even the existence of the fixed point is not evident. Proving global stability is also more complicated as it relies on induction on the component index. To prove the desired results, we use a new approach that employs bounds on the decay rate of the mean-field process and its monotonicity. We believe that this approach is more broadly applicable to other models where a fixed point does not admit a recursive relationship.

1.1 Related Works

In the last two decades, the Pod scheme has emerged as a widely used load balancing scheme due to its promising gains and minimal overhead. It has been studied extensively under various scaling limits and traffic conditions. The mean-field scaling limit for this scheme was first studied for exponential service time distributions in the seminal works [9, 10]. Their results were later generalised to general service time distributions in [21, 22]. The heavy traffic optimality of the Pod scheme has been established in [23, 24]. In [25], the analysis of the Pod scheme has been carried out for the case where the number of choices, d , is allowed to depend on the system size n and $d(n) = \omega(1)$. In this work, both the mean-field and Halfin-Whitt regimes are considered. In the mean-field limit, the Pod scheme has been shown to reach the same performance as the JSQ scheme. Recently, the Pod scheme has been analysed for different graph topologies. For example, in [26], the Pod scheme is analyzed for non-bipartite graphs and sufficient conditions on the graph sequence is obtained to match the result on complete graphs in the mean-field limit. For the bipartite graphs, the Pod scheme and its variants have been analysed in [27, 28]. In all cases, results similar to the complete graph setting have been obtained. For heterogeneous systems, the Pod scheme has been studied in [29, 30] where speed-aware versions of the Pod scheme have been shown to yield similar performance benefits.

The above mentioned results for the Pod scheme assume that on each arrival the dispatcher has the accurate knowledge of the queue lengths of the d sampled servers. However, this assumption may not be true in practice due to the issues discussed in the introduction. Recently, in [17], the *balls and bins problem* was studied under various noisy load comparison models. Here, n balls are placed into n bins sequentially and at each step a ball is placed into a bin from

a subset of d bins chosen at random. For the load-dependent error model discussed above, it has been shown that gap between the maximum and the average load is $O(\frac{g}{\log(g)} \log \log(n))$. This result motivates us to consider the effects of noisy load comparisons on the performance of the Po2 scheme in the dynamic setting where jobs are allowed to leave the system after being served. To the best of our knowledge, this is the first work that studies the Po2 scheme under the erroneous load comparison model in the dynamic setting.

2 SYSTEM MODEL

We consider a system consisting of n parallel servers, each with its own queue of infinite buffer size. Each server is able to process jobs at unit rate. Jobs arrive according to a Poisson process with a rate $n\lambda$ ($\lambda < 1$). Each job requires a random amount of work, independent and exponentially distributed with unit mean. The inter-arrival and job lengths are assumed to be independent of each other. A job dispatcher assigns each incoming job to a queue of a server where jobs are served according to the First-Come-First-Server (FCFS) scheduling discipline.

The job dispatcher uses the Po2 scheme to assign jobs to the servers¹. Under the classical Po2 scheme, a job is sent to the server with the minimum queue length among two servers, chosen uniformly at random. However, in practice, the server with the smaller queue length may not be always be correctly identified either due to outdated queue-length information at the dispatcher or due to an attacker misreporting the queue lengths sent from the servers to the dispatcher. Motivated by these scenarios, in this paper, we consider the following versions of the Po2 scheme where load comparisons are not always accurate. In the following, an error refers to the event where an arrival is sent to the server with the larger queue length among the two sampled servers.

2.1 Load-Dependent Error Model

In this model, an error occurs with probability $\epsilon \in [0, 1]$ only when the difference in queue lengths of the sampled servers is in the range $(0, g]$ for some constant $g \geq 0$. If the queue-length difference is strictly above g , then we assume that no error is made, i.e., the job is sent to the server with the smaller queue length. In case of a tie, we assume that an arbitrary tie breaking rule based on server indices is used. Without loss of generality (WLOG), we assume that servers are indexed from the index set $[n] = \{1, 2, \dots, n\}$, and, in case of a tie, the job is sent to the server with the smaller index among the two sampled servers. We refer to the Po2 scheme under this model of error as the Po2- (g, ϵ) scheme.

2.2 Load-Independent Error Model

In this model, an error occurs with probability $\epsilon \in [0, 1]$ independent of the current queue lengths of the sampled servers. More precisely, the incoming job is sent to the server having the higher queue length among the sampled servers with probability ϵ and with probability $1 - \epsilon$ it is sent to the server with the smaller queue length. Ties are broken in the same way as discussed before. For simplicity, we refer to the Po2 scheme under this model of error as the Po2- ϵ scheme.

¹We use Po2 scheme instead of the Po d scheme with $d \geq 3$ since the gain for $d \geq 3$ is marginal with respect to that for $d = 2$.

2.3 System State and Notations

To analyze the system under the schemes discussed above, we first introduce Markovian state descriptors of the system. We use two Markovian state descriptors. First, we define the queue-length vector at time $t \geq 0$ as $\mathbf{Q}^{(n)}(t) = (Q_k^{(n)}(t), k \in [n])$, where $Q_k^{(n)}(t)$ denotes the queue length of the k^{th} server. Second, we define the tail measure on the queue lengths at time t as $\mathbf{x}^{(n)}(t) = (x_i^{(n)}(t), i \geq 1)$, where $x_i^{(n)}(t)$ denotes the fraction of servers with at least i jobs at time t . For completeness, we set $x_i^{(n)}(t) = 1$ for all $i \leq 0$ and all $t \geq 0$. From the Poisson arrival and the exponential job size assumption it is clear that both $\mathbf{Q}^{(n)} = (\mathbf{Q}^{(n)}(t), t \geq 0)$ and $\mathbf{x}^{(n)} = (\mathbf{x}^{(n)}(t), t \geq 0)$ are Markov processes. When the system is stable, we denote by π_n the unique invariant measure of the process $\mathbf{x}^{(n)}$ and we use $\mathbf{x}^{(n)}(\infty)$ and $\mathbf{Q}^{(n)}(\infty)$ to denote the steady-state values of the processes $\mathbf{x}^{(n)}$ and $\mathbf{Q}^{(n)}$, respectively. As the load balancing scheme does not distinguish between servers, we have $\mathbb{P}(Q_i^{(n)}(t) \geq k) = \mathbb{E}[\mathbb{1}(Q_i^{(n)}(t) \geq k)] = (1/n) \sum_{i \in [n]} \mathbb{E}[\mathbb{1}(Q_i^{(n)}(t) \geq k)] = \mathbb{E}[x_k^{(n)}]$ for each $i \in [n]$ and each $t \in [0, \infty)$.

To define the state space of the process $\mathbf{x}^{(n)}$, we first define the space \bar{S} as $\bar{S} \triangleq \{\mathbf{s} = (s_i) : s_0 = 1, 1 \geq s_i \geq s_{i+1} \geq 0, \forall i \geq 1\}$. Note that the space \bar{S} is compact under the norm defined as $\|\mathbf{s}\| = \sup_{i \in \mathbb{Z}_+} \frac{|s_i|}{i+1}$ for $\mathbf{s} \in \bar{S}$. The process $\mathbf{Q}^{(n)}$ takes values in \mathbb{Z}_+^n and the process $\mathbf{x}^{(n)}$ takes values in the space $S^{(n)}$ defined as $S^{(n)} \triangleq \{\mathbf{s} \in \bar{S} : ns_i \in \mathbb{Z}_+ \forall i \geq 1\}$. We further define the space S as follows $S \triangleq \{\mathbf{s} \in \bar{S} : \|\mathbf{s}\|_1 < \infty\}$, where the ℓ_1 -norm, denoted by $\|\cdot\|_1$, is defined as $\|\mathbf{s}\|_1 \triangleq \sum_{i \geq 1} |s_i|$ for any $\mathbf{s} \in S$.

3 MAIN RESULTS AND INSIGHTS

In this section, we summarise our main results and discuss their consequences.

In the following theorem, we characterise the stability region for each of the two schemes discussed above.

Theorem 1 (Stability). (i) For any $g \geq 0, \epsilon \in [0, 1]$, and $n \geq 2$ the system under the Po2- (g, ϵ) scheme is stable (i.e., the process $\mathbf{x}^{(n)}$ is positive recurrent) if and only if $\lambda < 1$. Furthermore, for $\lambda < 1$, the steady-state average queue length per server is bounded above as

$$\mathbb{E}_{\pi_n}[Q_i^{(n)}(\infty)] = \mathbb{E}_{\pi_n} \left[\sum_{i \geq 1} x_i^n(\infty) \right] \leq \frac{(1 + g\mathbb{1}(\epsilon > 1/2))\lambda}{1 - \lambda}. \quad (1)$$

(ii) For any $\epsilon \in [0, 1]$ and $n \geq 2$, the system under the Po2- ϵ scheme is stable, if and only if $\lambda < \min\left(1, \frac{1}{2\epsilon}\right)$. Furthermore, for $\lambda < \min(1, 1/2\epsilon)$, the steady state average queue length per server is bounded above as

$$\mathbb{E}_{\pi_n}[Q_i^{(n)}(\infty)] = \mathbb{E}_{\pi_n} \left[\sum_{i \geq 1} x_i^n(\infty) \right] \leq \frac{\lambda}{1 - \max(1, 2\epsilon)\lambda}. \quad (2)$$

In addition to the stability regions, the above theorem gives uniform (in the system size n) bounds on the steady-state mean queue length per server for each scheme. These uniform bounds are crucial in establishing the tightness of stationary measures and justifying interchange of the limits in $\lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbf{x}^{(n)}(t) =$

$\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbf{x}^{(n)}(t)$, which shows that the mean-field approximation of the steady-state behaviour of the finite system is asymptotically exact.

The bounds in (1) and (2) also help us to compare the performance of the Po2- (g, ϵ) and the Po2- ϵ schemes to that of the random scheme. For example, when $\epsilon \leq 1/2$, both the upper bounds reduce to $\lambda/(1-\lambda)$ which is the steady-state average queue length per server under the random scheme. This implies that, under both models of error, the Po2 scheme performs better than the random scheme when the error probability $\epsilon \leq 1/2$. This is intuitive, as, for $\epsilon \leq 1/2$, an incoming job under the Po2 scheme is sent to the server with the smaller queue length more often than to the server with the larger queue length. For $\epsilon > 1/2$, however, the schemes may perform poorly in comparison to the random scheme (as both bounds become higher than $\lambda/(1-\lambda)$).




Figure 1: Mean response time of jobs under the Po2- (g, ϵ) scheme as a function of arrival rate λ for $\epsilon = 0.4, g = 100$.




Figure 2: Mean response time of jobs under the Po2- (g, ϵ) scheme and the random scheme as a function of arrival rate λ for $\epsilon = 0.8, g = 100$.

This is numerically verified in Figure 1 and Figure 2 for the Po2- (g, ϵ) scheme and in Figure 3 for the Po2- ϵ scheme. In each of these figures, we plot the steady-state mean response time of jobs as a function of the normalized arrival rate λ . From Figures 1 and 3, we observe that both the schemes outperform the random scheme when $\epsilon \leq 1/2$. For $\epsilon > 1/2$, however, the Po2- ϵ scheme becomes unstable for $\lambda \geq 1/2\epsilon$ and its performance becomes poorer than that of the random scheme for all $\lambda < 1$. For the Po2- (g, ϵ) scheme, we observe from Figure 2 that the system is stable for all $\lambda < 1$ even when $\epsilon > 1/2$. However, in this case, the performance of the Po2- (g, ϵ) scheme is poorer than that of the random scheme for small values of λ .

The usual approach of proving results similar to the ones stated in Theorem 1 consists of coupling and stochastic comparison with the




Figure 3: Comparison of the Po2- ϵ scheme for $\epsilon \in \{0.2, 0.4, 0.6, 0.8\}$ with the random scheme. We set $n = 200$.

random scheme. However, this approach does not work here since the random scheme can outperform each of the two schemes when $\epsilon > 1/2$. Instead, we use drifts of suitable Lyapunov functions to prove Theorem 1 which holds for all $\epsilon \in [0, 1]$. Establishing bounds on the drifts of Lyapunov functions is difficult for our schemes as only a subset of servers is compared at each arrival instant. We develop a generic approach through which the required bounds can be obtained for any scheme where queue lengths of multiple servers are compared to dispatch the incoming jobs.

The asymptotic (in n) results for the Po2- (g, ϵ) scheme are described in the following theorem.

Theorem 2. Po2- (g, ϵ)

(i) (Mean-Field Limit): Let $g \geq 0, \epsilon \in [0, 1]$ and assume $\mathbf{x}^{(n)}(0) \in S^{(n)}$ for each n and $\mathbf{x}^{(n)}(0) \xrightarrow{a.s.} \mathbf{u} \in S$ under ℓ_1 as $n \rightarrow \infty$. Then, for each $T \geq 0$, we have $\sup_{t \in [0, T]} \|\mathbf{x}^{(n)}(t) - \mathbf{x}(t)\|_1 \xrightarrow{a.s.} 0$ where $\mathbf{x} = (\mathbf{x}(t) = (x_i(t), i \geq 1), t \geq 0)$ satisfies $\mathbf{x}(0) = \mathbf{u}$ and for $t \geq 0$ and $i \geq 1$

$$\dot{x}_i(t) = G_i(\mathbf{x}(t)) \triangleq \lambda p_{i-1}(\mathbf{x}(t)) - (x_i(t) - x_{i+1}(t)). \quad (3)$$

Here for each $i \geq 1, G_i$ is the i^{th} component of the function $\mathbf{G} = (G_i, i \geq 1) : S \rightarrow \mathbb{R}^\infty$ and, for $\mathbf{s} \in S, p_{i-1}(\mathbf{s})$ is defined as

$$p_{i-1}(\mathbf{s}) = (s_{i-1} - s_i) [2\epsilon(s_{i+g} + s_{i-1-g}) + (1 - 2\epsilon)(s_i + s_{i-1})]. \quad (4)$$

(ii) (Mean-Field Steady State Behaviour): For $g \geq 0, \epsilon \in [0, 1]$, and $\lambda < 1$, there exists a unique $\mathbf{x}^* \in S$ such that, $\mathbf{G}(\mathbf{x}^*) = \mathbf{0}$. In addition, \mathbf{x}^* satisfies

$$x_1^* = \lambda, \quad (5)$$

$$x_k^* = \lambda \left[2\epsilon \left(x_{k-1}^* x_{k-1-g}^* - \sum_{i=k}^{k+g-1} x_i^* (x_{i-1-g}^* - x_{i-g}^*) \right) + (1 - 2\epsilon)(x_{k-1}^*)^2 \right], \text{ for } k \geq 2. \quad (6)$$

Moreover, any solution $\mathbf{x}(t)$ of (3) with $\mathbf{x}(0) \in S$ converges to \mathbf{x}^* in ℓ_1 as $t \rightarrow \infty$. Therefore, the sequence $(\pi_n)_{n \geq 1}$ converges weakly to the Dirac measure $\delta_{\mathbf{x}^*}$ concentrated on \mathbf{x}^* as $n \rightarrow \infty$.

(iii) (Heavy-Traffic Limit): Furthermore, in the heavy traffic as $\lambda \rightarrow 1$, we have

$$\limsup_{\lambda \rightarrow 1} \frac{T_2^{g, \epsilon}(\lambda)}{\log(T_1(\lambda))} \leq \frac{g+1}{\log(2)}, \quad (7)$$

where $T_2^{g, \epsilon}(\lambda) = (1/\lambda) \sum_{i \geq 1} x_i^*$ is the steady state limiting (as $n \rightarrow \infty$) average response time of jobs under the Po2- (g, ϵ) scheme and

$T_1(\lambda) = 1/(1 - \lambda)$ is the steady state average response time of jobs under the random scheme.

Hence, the process \mathbf{x} , defined in Theorem 2.(i), characterises the dynamics of the system in the limit as $n \rightarrow \infty$. This will be referred to as the *mean-field limit* of the system or the *mean-field process*. The evolution of \mathbf{x} , described in (3), can be explained as follows. For the n^{th} system, the component $x_i^n(t)$ increases by $1/n$ when a job joins a server with queue length exactly $i - 1$. The rate at which this happens is $n\lambda p_{i-1}(\mathbf{x})$, where $p_{i-1}(\mathbf{s})$, for $\mathbf{s} \in S$, is the probability that an arrival joins a server with queue-length $i - 1$ when the system is in state \mathbf{s} . The expression of $p_{i-1}(\mathbf{s})$ in (4) can be obtained as follows. Under the Po2- (g, ϵ) scheme, a job joins a server with queue length $i - 1$ under the following scenarios: (1) One of the two sampled servers is of queue length exactly $i - 1$ and the other sampled server is of queue length is at least $i + g$. This occurs with probability $2(s_{i-1} - s_i)s_{i+g}$ and, in this case, the job joins the server queue length $i - 1$ with probability 1. (2) One of the sampled server is of queue length $i - 1$ and the other sampled server's queue length lies in the range $\{i - 1 - g, \dots, i - 2\}$. This occurs with probability $2(s_{i-1} - s_i)(s_{i-1-g} - s_{i-1})$, and, in this case the server with queue length $i - 1$ is selected with probability ϵ . (3) One of the sampled server is of queue length $i - 1$ and the other sampled server's queue length lies in the range $\{i + 1, \dots, i - 1 + g\}$. This occurs with probability $2(s_{i-1} - s_i)(s_i - s_{i+g})$ and, in this case, the server with queue length $i - 1$ is selected with probability $(1 - \epsilon)$. (4) Finally, both the sampled servers can have the same queue length $i - 1$ with probability $(s_{i-1} - s_i)^2$, and, in this case, the job joins a server with queue length $i - 1$ with probability 1. Combining the above probabilities, we obtain the expression for $p_{i-1}(\mathbf{s})$. Similarly, the component $x_i^n(t)$ decreases by $1/n$ when a job leaves a server with queue length i and this occurs with rate $n(x_i - x_{i+1})$. Hence, the total expected rate of change (drift) in the component $x_i^n(t)$ is given by $G_i(\mathbf{x}^n(t)) = p_{i-1}(\mathbf{x}^n(t)) - (x_i^n(t) - x_{i+1}^n(t))$. In the limit as $n \rightarrow \infty$, this becomes the rate of change of $x_i(t)$.

In part (ii) of Theorem 2, we show that, as $t \rightarrow \infty$, the mean-field process \mathbf{x} converges in ℓ_1 to the unique point $\mathbf{x}^* \in S$ at which $\mathbf{G}(\mathbf{x}^*) = \mathbf{0}$. This point \mathbf{x}^* is referred to as the *fixed point* of the mean-field since starting at this point the mean-field remains at this point at all times. Since by part (i) we have $\mathbf{x}^n(t) \rightarrow \mathbf{x}(t)$ almost surely for each $t \geq 0$, the convergence to the fixed point implies $\lim_{t \rightarrow \infty} \lim_{n \rightarrow \infty} \mathbf{x}^n(t) = \lim_{n \rightarrow \infty} \lim_{t \rightarrow \infty} \mathbf{x}^n(t) = \mathbf{x}^*$, which, in turn, means that the fixed point \mathbf{x}^* characterises the steady-state behaviour of the system in the limit as $n \rightarrow \infty$. In particular, $\lim_{n \rightarrow \infty} \mathbb{P}(Q_i^{(n)}(\infty) \geq k) = x_k^*$.

In Theorem 2.(iii) we compare the steady-state mean response time of jobs under the Po2- (g, ϵ) scheme to that under the random scheme when the traffic is high (i.e., $\lambda \rightarrow 1$). Note that (7) implies that $T_2^{g, \epsilon}(\lambda) = O(\log T_1(\lambda))$ as $\lambda \rightarrow 1$. Furthermore, by the previous part of the theorem, the steady-state mean response time of the jobs under the Po2- (g, ϵ) scheme converges as $n \rightarrow \infty$ to $T_2^{g, \epsilon}(\lambda)$. Hence, this result shows that, when the system is heavily loaded, the mean response time of jobs under the Po2- (g, ϵ) scheme is exponentially smaller than that under the random scheme. This is also verified in Figure 2 for $\epsilon = 0.8$ and $g = 100$. Note that for such high error rates, the mean response time of jobs under the Po2- (g, ϵ) policy can be larger than that under the random scheme for low values of

λ . However, when λ is close to its maximum value 1, the Po2- (g, ϵ) scheme performs exponentially better than the random scheme for all values of g and ϵ . This implies that the advantage of having an additional choice in the Po2 scheme outweighs the negative impact the comparison errors when the traffic is high.

The main difficulty in proving Theorem 2 is that the fixed point \mathbf{x}^* cannot be found in closed form. This is because each component x_k^* in (6) depends not only on the previous components but also on the next g components. This makes it hard to characterise the fixed point; indeed, even the existence of such \mathbf{x}^* in S is not evident. This also makes proving the global stability difficult as it uses induction on the component index k . To overcome these difficulties, we use the monotonicity of the mean-field and uniform bounds on its tails. We believe that this new approach is generally applicable to similar systems where the fixed point cannot be found in closed form.

We now present the asymptotic results for the Po2- ϵ scheme in the following theorem.

Theorem 3. Po2- ϵ

(i) (*Mean-Field Limit*): Let $\epsilon \in [0, 1]$ and assume $\mathbf{x}^{(n)}(0) \in S^{(n)}$ for each n and $\mathbf{x}^{(n)}(0) \xrightarrow{a.s.} \mathbf{u} \in S$ under ℓ_1 as $n \rightarrow \infty$. Then, for each $T \geq 0$, we have $\sup_{t \in [0, T]} \|\mathbf{x}^n(t) - \mathbf{x}(t)\|_1 \xrightarrow{a.s.} 0$ where $\mathbf{x} = (\mathbf{x}(t) = (x_i(t), i \geq 1), t \geq 0)$ satisfies $\mathbf{x}(0) = \mathbf{u}$ and for $t \geq 0$ and $i \geq 1$

$$\dot{x}_i(t) = \lambda p_{i-1}(\mathbf{x}(t)) - (x_i(t) - x_{i+1}(t)), \quad (8)$$

where

$$p_{i-1}(\mathbf{s}) = (1 - \epsilon)(s_{i-1}^2 - s_i^2) + \epsilon((1 - s_i)^2 - (1 - s_{i-1})^2), \quad \mathbf{s} \in S. \quad (9)$$

We refer to the process \mathbf{x} as the *mean-field limit* of the sequence $(\mathbf{x}^{(n)})_{n \geq 1}$.

(ii) (*Mean-Field Steady State Behaviour*): For $\epsilon \leq 1/2$ and $\lambda < 1$, there exists $\mathbf{x}^* \in S$ such that, if $\mathbf{x}(0) = \mathbf{x}^*$, then $\mathbf{x}(t) = \mathbf{x}^*$ for all $t \geq 0$. Furthermore, \mathbf{x}^* satisfies the following recursion

$$x_1^* = \lambda, \quad x_i^* = \lambda \left[(1 - 2\epsilon)(x_{i-1}^*)^2 + 2\epsilon x_{i-1}^* \right] \quad \forall i \geq 2. \quad (10)$$

Moreover, any solution $\mathbf{x}(t)$ of (8) with $\mathbf{x}(0) \in S$ converges to \mathbf{x}^* in ℓ_1 , i.e., $\|\mathbf{x}(t) - \mathbf{x}^*\|_1 \rightarrow 0$ as $t \rightarrow \infty$. The above results imply that the sequence $(\pi_n)_n$ converges weakly to the Dirac measure $\pi_* = \delta_{\mathbf{x}^*}$ concentrated on \mathbf{x}^* as $n \rightarrow \infty$.

(iii) (*Heavy-Traffic Limit*): For $\epsilon \leq 1/2$ and $\lambda < 1$, we have

$$\lim_{\lambda \rightarrow 1^-} \frac{T_2^\epsilon(\lambda)}{\log(T_1(\lambda))} = \frac{1}{\log(2 - 2\epsilon)}, \quad (11)$$

where $T_2^\epsilon(\lambda) = \frac{\sum_{k=1}^{\infty} x_k^*}{\lambda}$ is the limiting (as $n \rightarrow \infty$) steady state average response time of jobs under the Po2- ϵ scheme.

In parts (i) and (ii) of Theorem 3, we characterize the mean-field limit \mathbf{x} and its fixed point \mathbf{x}^* under the Po2- ϵ scheme. As before, we show that the fixed point is unique and globally asymptotically stable. In the last part (part iii) of Theorem 3, we compare the mean response time of jobs under the Po2- ϵ scheme to that under the random scheme in the limit as $n \rightarrow \infty$. Our result indicates that when n is large and λ is close to 1, the steady state mean response time of jobs under the Po2- ϵ scheme satisfies $T_2^\epsilon(\lambda) \approx c_\epsilon \log(T_1(\lambda))$, where $c_\epsilon = 1/(\log(2 - 2\epsilon))$. This means that an exponential reduction in the steady state mean response time is achieved as long as $\epsilon \leq 1/2$. Hence, the Po2- ϵ scheme retains the benefits of the Po2 scheme as long as $\epsilon \leq 1/2$.

The proof of Theorem 3 is much simpler than that of Theorem 2 as, for the Po2- ϵ scheme, a recursive relationship among the components x_k^* exists (as given by (10)). Hence, the proofs for the Po2- ϵ scheme follow standard techniques. We omit them to conserve space. Interested readers can follow the extended version [31] of this paper where the proofs can be found.

4 STABILITY AND UNIFORM BOUNDS

In this section, we find the stability regions for the Po2- (g, ϵ) and the Po2- ϵ schemes and derive uniform bounds on the steady-state queue length per server (Theorem 1) using drifts of appropriate Lyapunov functions. We first develop a general framework to analyse any load balancing scheme that compares the queue lengths of two uniformly sampled servers to dispatch every job. Note that it is easy to generalise this framework further to cases where more than two servers are sampled and the sampling is not necessarily uniform.

For any function $V : \mathbb{Z}_+^n \rightarrow [0, \infty)$, the drift of $D_{\mathbf{Q}^n} V$ is defined as the expected rate of change in the value of the function along the trajectory of the process \mathbf{Q}^n given the current state. More precisely,

$$\begin{aligned} D_{\mathbf{Q}^n} V(\mathbf{Q}) &= \lim_{h \rightarrow 0} \frac{1}{h} \mathbb{E}[V(\mathbf{Q}^{(n)}(t+h)) - V(\mathbf{Q}^{(n)}(t)) | \mathbf{Q}^n(t) = \mathbf{Q}] \\ &= \sum_{i=1}^n [r_i^{+,n}(\mathbf{Q})(V(\mathbf{Q} + \mathbf{e}_i^{(n)}) - V(\mathbf{Q})) \\ &\quad + r_i^{-,n}(\mathbf{Q})(V(\mathbf{Q} - \mathbf{e}_i^{(n)}) - V(\mathbf{Q}))], \end{aligned} \quad (12)$$

where $\mathbf{e}_i^{(n)}$ denotes the n -dimensional unit vector with one in the i^{th} position; $r_i^{\pm, n}(\mathbf{Q})$ are the transition rates from the state \mathbf{Q} to the states $\mathbf{Q} \pm \mathbf{e}_i^{(n)}$. According to the Foster-Lyapunov theorem (Proposition D.3 of [32]), to prove the stability or positive recurrence of the process $\mathbf{Q}^{(n)}$, it is sufficient to show the existence of at least one function $V : \mathbb{Z}_+^n \rightarrow [0, \infty)$ such that $V(\mathbf{Q}) \rightarrow \infty$ when $\|\mathbf{Q}\|_1 \rightarrow \infty$ and $D_{\mathbf{Q}^n} V(\mathbf{Q}) < 0$ for all states \mathbf{Q} lying outside a compact subset of the state-space. To further obtain uniform bounds on the stationary queue lengths, we use the fact (Proposition 1 of [33]) that $\mathbb{E}_{\pi_n}[D_{\mathbf{Q}^n} V(\mathbf{Q}(\infty))] \geq 0$ if $D_{\mathbf{Q}^n} V(\mathbf{Q})$ is uniformly bounded for all states $\mathbf{Q} \in \mathbb{Z}_+^n$.

The rate of departure from the i^{th} server is given by $r_i^{-, n}(\mathbf{Q}) = \mathbb{1}(Q_i > 0)$.

For any scheme which compares the states of two servers to dispatch the job to one of the servers, we define the class of an arrival as the (unordered) pair (i, j) of servers sampled at the arrival instant. Let \mathcal{C} denote the collection of all such classes. Since $|\mathcal{C}| = \binom{n}{2}$ and a job is equally likely to belong to one of these classes, the arrival rate of any class $(i, j) \in \mathcal{C}$ is $n\lambda / \binom{n}{2} = 2\lambda / (n-1)$. Hence, we can write the rate of arrival to the i^{th} server as

$$r_i^{+, n}(\mathbf{Q}) = \frac{2\lambda}{n-1} \sum_{j \in [n], j \neq i} p(Q_i, Q_j), \quad (13)$$

where $p(Q_i, Q_j)$ is the probability that a class (i, j) job joins the server i when the queue lengths of servers i and j are Q_i and Q_j , respectively. Note that the probability $p(Q_i, Q_j)$ depends on the load balancing scheme used by the dispatcher. The exact expression of $p(Q_i, Q_j)$ for each scheme is given later, but it is important to

note that $p(Q_i, Q_j) + p(Q_j, Q_i) = 1$ since a class (i, j) job joins either server i or server j with probability 1.

Now, for the Lyapunov function $V : \mathbb{Z}_+^n \rightarrow [0, \infty)$ defined as $V(\mathbf{Q}) = \sum_{i=1}^n Q_i^2$, the drift given in (12) simplifies to

$$D_{\mathbf{Q}^n} V(\mathbf{Q}) = \sum_{i=1}^n \left\{ 2[r_i^{+, n}(\mathbf{Q}) - r_k^{-, n}(\mathbf{Q})]Q_i + [r_i^{+, n}(\mathbf{Q}) + r_i^{-, n}(\mathbf{Q})] \right\},$$

which upon further simplification gives

$$D_{\mathbf{Q}^n} V(\mathbf{Q}) = 2 \sum_{i=1}^n r_i^{+, n}(\mathbf{Q})Q_i - 2 \sum_{i \in [n]} Q_i + n\lambda + B(\mathbf{Q}), \quad (14)$$

where $B(\mathbf{Q}) = \sum_{i \in [n]} \mathbb{1}(Q_i > 0)$ represents the number of busy servers when system is in state \mathbf{Q} . In the above, we have used the facts $r_i^{-, n}(\mathbf{Q}) = \mathbb{1}(Q_i > 0)$ and $\sum_{i \in [n]} r_i^{+, n}(\mathbf{Q}) = n\lambda$. Moreover, using (13), the first term in the expression of the drift can be written as

$$\begin{aligned} \sum_{i=1}^n r_i^{+, n}(\mathbf{Q})Q_i &= \frac{2\lambda}{(n-1)} \sum_{i \in [n]} \sum_{j \in [n], j \neq i} Q_i p(Q_i, Q_j) \\ &= \frac{2\lambda}{(n-1)} \sum_{(i,j) \in \mathcal{C}} (Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i)). \end{aligned} \quad (15)$$

Thus, to obtain the stability region and the uniform bound on steady-state queue length, we need to obtain upper bounds on $Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i)$ for each scheme.

4.1 Po2- (g, ϵ) Scheme

For the Po2- (g, ϵ) scheme, the probability $p(Q_i, Q_j)$ for any class $(i, j) \in \mathcal{C}$ is given by

$$\begin{aligned} p(Q_i, Q_j) &= \mathbb{1}(Q_j - Q_i \geq g+1) + (1-\epsilon)\mathbb{1}(Q_j - Q_i \in (0, g]) \\ &\quad + \epsilon\mathbb{1}(Q_i - Q_j \in (0, g]) + \mathbb{1}(Q_i = Q_j, i < j). \end{aligned} \quad (16)$$

Using the above expression, we obtain the following bound for Po2- (g, ϵ) scheme.

Lemma 4. For $g \geq 0$, $\epsilon \in [0, 1]$, and for any class $(i, j) \in \mathcal{C}$, under the Po2- (g, ϵ) , scheme we have

$$Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) \leq \frac{Q_i + Q_j}{2} + g\mathbb{1}(\epsilon > 1/2). \quad (17)$$

PROOF. To prove the lemma, we first observe that for any $a \leq b$, and $w_1 = 1 - w_2 \in [1/2, 1]$ we have

$$w_1 a + w_2 b \leq \frac{a+b}{2}. \quad (18)$$

The result of the lemma is direct when $Q_i = Q_j$. So, we consider the case $Q_i < Q_j$. Note that the proof for $Q_i > Q_j$ is exactly the same with Q_i and Q_j interchanged. For $\epsilon \leq 1/2$, using (16) and $Q_i < Q_j$, we have $p(Q_i, Q_j) = \mathbb{1}(Q_j - Q_i \geq g+1) + (1-\epsilon)\mathbb{1}(Q_j - Q_i \in (0, g]) \in [1/2, 1]$. Since $p(Q_i, Q_j) = 1 - p(Q_j, Q_i)$, using (18), we have $Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) \leq (Q_i + Q_j)/2$. Next, we note from (16) and $Q_i < Q_j$ that

$$\begin{aligned} Q_i p(Q_i, Q_j) &= Q_i (\mathbb{1}(Q_j - Q_i \geq g+1) + (1-\epsilon)\mathbb{1}(Q_j - Q_i \in (0, g])) \\ &= Q_i ((1-\epsilon) + \epsilon\mathbb{1}(Q_j - Q_i \geq g+1)), \end{aligned}$$

which gives

$$Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) = Q_i ((1 - \epsilon) + \epsilon \mathbb{1}(Q_j - Q_i \geq g + 1)) \\ + Q_j \epsilon \mathbb{1}(Q_j - Q_i \in (0, g]).$$

Therefore, using $Q_j \mathbb{1}(Q_j - Q_i \in (0, g]) \leq (g + Q_i) \mathbb{1}(Q_j - Q_i \in (0, g]) \leq (g + (Q_i + Q_j)/2) \mathbb{1}(Q_j - Q_i \in (0, g])$ and $Q_i < (Q_i + Q_j)/2 + g$ in the above, we obtain

$$Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) \leq g + \frac{Q_i + Q_j}{2}.$$

This completes the proof. \square

Proof of Theorem 1.(i): Using the bound of Lemma 4, the RHS of (15) can be bounded as

$$\sum_{i=1}^n r_i^{+,n}(\mathbf{Q}) Q_i \leq \frac{2\lambda}{(n-1)} \sum_{(i,j) \in C} \left\{ \frac{Q_i + Q_j}{2} + g \mathbb{1}(\epsilon > 1/2) \right\} \\ = \frac{2\lambda}{(n-1)} \left(\frac{(n-1)}{2} \sum_{i \in [n]} Q_i + g \mathbb{1}(\epsilon > 1/2) \binom{n}{2} \right) \\ = \lambda \sum_{i \in [n]} Q_i + n\lambda g \mathbb{1}(\epsilon > 1/2) \quad (19)$$

Therefore, using (19) in (14), we can upper bound the drift $D_{\mathbf{Q}^{(n)}} V(\mathbf{Q})$ for the Po2- (g, ϵ) scheme as

$$D_{\mathbf{Q}^{(n)}} V(\mathbf{Q}) \leq 2(\lambda - 1) \sum_{i=1}^n Q_i + 2n\lambda g \mathbb{1}\left(\epsilon > \frac{1}{2}\right) + n\lambda + B(\mathbf{Q}). \quad (20)$$

Now, since $B(\mathbf{Q}) \leq n$ and $\lambda < 1$, the drift $D_{\mathbf{Q}^{(n)}} V(\mathbf{Q})$ is strictly negative whenever $\sum_{i \in [n]} Q_i > n(\lambda(2g \mathbb{1}(\epsilon > 1/2) + 1) + 1)/(2(1 - \lambda))$, and is bounded above by $n(\lambda(2g \mathbb{1}(\epsilon > 1/2) + 1) + 1)$, otherwise. This shows that the system under the Po2- (g, ϵ) scheme is stable for all $\lambda < 1$. The necessity of this condition for stability can be established easily by showing that the drift of the Lyapunov function $V_1(\mathbf{Q}) = \sum_{i \in [n]} Q_i$ is always non-negative when $\lambda \geq 1$.

To prove (1), recall from the previous paragraph that $D_{\mathbf{Q}^{(n)}} V(\mathbf{Q})$ is uniformly bounded by $n(\lambda(2g \mathbb{1}(\epsilon > 1/2) + 1) + 1)$ for all $\mathbf{Q} \in \mathbb{Z}_+^n$. This implies that $\mathbb{E}_{\pi_n} [D_{\mathbf{Q}^{(n)}} V(\mathbf{Q}^{(n)}(\infty))] \geq 0$. Therefore, taking expectation of (20) and using the rate conservation equation $\mathbb{E}_{\pi_n} [B(\mathbf{Q})] = n\lambda$ (which holds in steady-state), we obtain

$$\frac{n\lambda(1 + g \mathbb{1}(\epsilon > 1/2))}{(1 - \lambda)} \geq \mathbb{E}_{\pi_n} \left[\sum_{i \in [n]} Q_i^{(n)}(\infty) \right] = n \mathbb{E}_{\pi_n} [Q_i^{(n)}(\infty)],$$

where last equality follows due to the exchangeability of π_n . \square

4.2 Po2- ϵ Scheme

For the Po2- ϵ scheme, $p(Q_i, Q_j)$ for any class $(i, j) \in C$ is given by

$$p(Q_i, Q_j) = (1 - \epsilon) \mathbb{1}(Q_i < Q_j) + \epsilon \mathbb{1}(Q_i > Q_j) + \mathbb{1}(Q_i = Q_j, i < j). \quad (21)$$

Using the expression above, we obtain the following bound.

Lemma 5. For $\epsilon \in [0, 1]$ and for any class $(i, j) \in C$, under the Po2- ϵ , scheme we have

$$Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) \leq (Q_i + Q_j) \max(1/2, \epsilon). \quad (22)$$

PROOF. For $Q_i = Q_j$ the above inequality follows directly. Similar to the proof of Lemma 4, it is sufficient to consider the case $Q_i < Q_j$. For $\epsilon \leq 1/2$, using (21) we have $p(Q_i, Q_j) = (1 - \epsilon) \in [1/2, 1]$. Therefore, from (18), it follows that $Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) \leq (Q_i + Q_j)/2$. When $\epsilon > 1/2$, for $Q_i < Q_j$, we can write

$$Q_i p(Q_i, Q_j) + Q_j p(Q_j, Q_i) = Q_i(1 - \epsilon) + Q_j \epsilon \leq \epsilon(Q_i + Q_j).$$

Hence, the proof is complete. \square

Proof of Theorem 1.(ii): Using the bound of Lemma 5, the RHS of (15) can be bounded as

$$\sum_{i=1}^n r_i^{+,n}(\mathbf{Q}) Q_i \leq \frac{2\lambda}{(n-1)} \sum_{(i,j) \in C} (Q_i + Q_j) \max(1/2, \epsilon), \\ = \lambda \max(1, 2\epsilon) \sum_{i \in [n]} Q_i. \quad (23)$$

Therefore, using (23) in (14), we upper-bound the drift $D_{\mathbf{Q}^{(n)}} V(\mathbf{Q})$ for the Po2- ϵ scheme as

$$D_{\mathbf{Q}^{(n)}} V(\mathbf{Q}) \leq 2(\lambda \max(1, 2\epsilon) - 1) \sum_{i \in [n]} Q_i + n\lambda + B(\mathbf{Q}). \quad (24)$$

Since $B(\mathbf{Q}) \leq n$, the above implies that, for $\lambda < \min(1, 1/2\epsilon)$, the drift is strictly negative whenever $\sum_{i \in [n]} Q_i > n(\lambda + 1)/(1 - \lambda \max(1, 2\epsilon))$. This shows that the system under Po2- ϵ scheme is stable for all $\lambda < \min(1, 1/2\epsilon)$. Furthermore, since $\sup_{\mathbf{Q} \in \mathbb{Z}_+^n} D_{\mathbf{Q}^{(n)}} V(\mathbf{Q}) \leq n(\lambda + 1)$, using $\mathbb{E}_{\pi_n} [D_{\mathbf{Q}^{(n)}} V(\mathbf{Q}^{(n)}(\infty))] \geq 0$, we obtain

$$\frac{n\lambda}{(1 - \lambda \max(1, 2\epsilon))} \geq \mathbb{E}_{\pi_n} \left[\sum_{i \in [n]} Q_i^{(n)}(\infty) \right] = n \mathbb{E}_{\pi_n} [Q_i^{(n)}(\infty)],$$

which proves (2).

Next we prove that for $\lambda \geq \min(1, \frac{1}{2\epsilon})$ the system is unstable. For $\epsilon \leq 1/2$ and $\lambda \geq 1$, the process $\mathbf{Q}^{(n)}$ is not positive recurrent. This follows using the same argument as used in the stability proof of the Po2- (g, ϵ) scheme. Now, for $\epsilon > 1/2$ and $2\lambda\epsilon > 1$, we consider the Lyapunov function $V_2(\mathbf{Q}) = Q_{i^*}(\mathbf{Q})$, where $i^*(\mathbf{Q}) = \arg \max_{i \in [n]} Q_i$ and $i^*(\mathbf{Q})$ is the minimum such index. Using (12), the drift of the function $V_2(\mathbf{Q})$ can be written as

$$D_{\mathbf{Q}^{(n)}} V_2(\mathbf{Q}) = \left(r_{i^*}^{+,n}(\mathbf{Q}) - \mathbb{1}(Q_{i^*}(\mathbf{Q}) > 0) \right) \geq \left(r_{i^*}^{+,n}(\mathbf{Q}) - 1 \right).$$

From (13), we have

$$r_{i^*}^{+,n}(\mathbf{Q}) = \frac{2\lambda}{n-1} \sum_{j \neq i^*(\mathbf{Q})} p(Q_i, Q_j) \geq \frac{2\lambda}{n-1} \sum_{j \neq i^*(\mathbf{Q})} \epsilon = 2\epsilon\lambda.$$

where the inequality follows from (21) since $Q_{i^*}(\mathbf{Q}) \geq Q_j$ for any $j \in [n]$ and $i^*(\mathbf{Q})$ is the minimum such index. Hence, $D_{\mathbf{Q}^{(n)}} V_2(\mathbf{Q}) \geq 2\epsilon\lambda - 1 \geq 0$ for all $\mathbf{Q} \in \mathbb{Z}_+^n$. Furthermore, since $D_{\mathbf{Q}^{(n)}} V_2(\mathbf{Q}) \leq n\lambda$, the result follows from the Foster-Lyapunov criterion for transience and null recurrence (Theorem 3.3.10 of [34]). \square

5 MEAN-FIELD ANALYSIS OF THE PO2- (g, ϵ) SCHEME

In this section, we prove the main results for the Po2- (g, ϵ) scheme stated in Theorem 2.

5.1 Mean-Field Limit of the Po2-(g, ϵ) Scheme

First, we establish the mean-field limit of the Po2-(g, ϵ) scheme (Theorem 2.(i)) for any $g \geq 0$ and $\epsilon \in [0, 1]$. Note that under the Po2-(g, ϵ) scheme, the rate of transition of the process $\mathbf{x}^{(n)}$ from state $\mathbf{x} \in S^{(n)}$ to state $\mathbf{y} \in S^{(n)}$ is given by

$$r_{\mathbf{x}, \mathbf{y}}^{(n)} = \begin{cases} n\lambda p_{i-1}(\mathbf{x}), & \text{if } \mathbf{y} = \mathbf{x} + \mathbf{e}_i/n, \forall i \geq 1, \\ n(x_i - x_{i+1}), & \text{if } \mathbf{y} = \mathbf{x} - \mathbf{e}_i/n \end{cases}, \quad (25)$$

where $p_{i-1}(\mathbf{x})$ is as defined in (4) and \mathbf{e}_i is the i^{th} unit vector in \mathbb{R}^∞ . Clearly, the above rates satisfy the transition structure for a *density-dependent jump Markov chain* [9, 35]. Furthermore, it is easy to verify that $\sum_{\mathbf{y} \in S} r_{\mathbf{x}, \mathbf{y}}^{(n)} < n(\lambda + 1)$ for all $\mathbf{x} \in S$ and the function $\mathbf{G} : S \rightarrow \mathbb{R}^\infty$ is Lipschitz under the ℓ_1 -norm with a Lipschitz constant of $L_\lambda^\epsilon = \lambda(16\epsilon + 4) + 2$. Hence, using the Kurtz's theorem for density-dependent jump Markov processes [[36], Chapter 8], we obtain the desired result.

5.2 Mean-Field Steady State Behaviour for the Po2-(g, ϵ) Scheme

We now turn to the proof of Theorem 2.(ii) which shows that the mean-field process \mathbf{x} given by (3) has a unique fixed point \mathbf{x}^* which satisfies (5) and (6). Moreover, we show that the fixed point \mathbf{x}^* is globally stable, i.e., all trajectories of the mean-field process \mathbf{x} starting in S converges to \mathbf{x}^* .

For any $\mathbf{u} \in S$, let $\mathbf{x}(t, \mathbf{u})$ denote the trajectory of the mean-field process starting at state \mathbf{u} . Further, define $v_k(t, \mathbf{u}) = \sum_{i \geq k} x_i(t, \mathbf{u})$ and $v_k(\mathbf{u}) = \sum_{i \geq k} u_i$ for each $k \geq 1$. When the context is clear, we shall drop the dependence of the trajectory on the initial state \mathbf{u} and on the time t .

Lemma 6. *Let $g \geq 0, \epsilon \in [0, 1]$. The following statements hold for the process \mathbf{x} defined in Theorem 2.*

- (1) *If $\mathbf{u} \in S$, then, for any $t \geq 0$, we have $\mathbf{x}(t, \mathbf{u}) \in S$.*
- (2) *For any $\mathbf{u}, \mathbf{u}' \in S$ satisfying $\mathbf{u} \leq \mathbf{u}'$ we have $\mathbf{x}(t, \mathbf{u}) \leq \mathbf{x}(t, \mathbf{u}')$ for all $t \geq 0$, where the inequality \leq is component-wise.*

PROOF. We first note from (3) and (4) that

$$\begin{aligned} \dot{x}_i &\leq \lambda p_{i-1}(\mathbf{x}) \leq \lambda [2\epsilon(x_{i-1-g}^2 - x_i^2) + (1-2\epsilon)(x_{i-1}^2 - x_i^2)] \\ &\leq \lambda \max(1, 2\epsilon)x_{i-1-g}^2 \leq M_{\lambda, \epsilon} x_{i-1-g}, \end{aligned}$$

where $M_{\lambda, \epsilon} = \max(1, 2\epsilon)$. This implies that for each $i \geq 1$ we have

$$x_i(t) \leq x_i(0) + M_{\lambda, \epsilon} \int_0^t x_{i-1-g}(s) ds.$$

Using the above recursively for each $i \geq 1$, we obtain

$$x_i(t) \leq x_i(0) + \sum_{k=0}^{i-g-1} x_k(0) \frac{(M_{\lambda, \epsilon} t)^{i-g-k}}{(i-g-k)!}.$$

Summing the above for all $i \geq 1$ we obtain

$$v_1(\mathbf{x}(t)) \leq (g+1 + v_1(\mathbf{x}(0))) \exp(M_{\lambda, \epsilon} t).$$

This shows that if $v_1(\mathbf{x}(0)) < \infty$, then $v_1(\mathbf{x}(t)) < \infty$ for all t , thus establishing the first part of the lemma.

For the second part, from (3) and (4) we note that

$$\frac{\partial G_i(\mathbf{x})}{\partial x_{i-1}} = \lambda [2((1-\epsilon)x_{i-1} + x_{i+g}) + 2\epsilon(x_{i-1-g} - x_{i-1})] \geq 0.$$

It is also clear that $\frac{\partial G_i(\mathbf{x})}{\partial x_{i+g}}$ and $\frac{\partial G_i(\mathbf{x})}{\partial x_{i-g-1}}$ and $\frac{\partial G_i(\mathbf{x})}{\partial x_{i+1}}$ are non-negative. Hence, we have for each $i \geq 1$, $G_i(\mathbf{x}) = \lambda p_{i-1}(\mathbf{x}) - (x_i - x_{i+1})$ is non-decreasing with respect to all $x_k, k \neq i$. Now the result follows from Theorem 5.3 of [37]. \square

The second property stated in Lemma 6 is called the *quasi-monotonicity* of the process \mathbf{x} . This property ensures that if the mean-field process starts from the idle initial state, i.e., if $\mathbf{x}(0) = \mathbf{e}_0 = (1, 0, 0, \dots)$, then it is monotonically non-decreasing in time, i.e.,

$$\mathbf{x}(t_1, \mathbf{e}_0) \leq \mathbf{x}(t_2, \mathbf{e}_0), \quad 0 \leq t_1 \leq t_2 < \infty. \quad (26)$$

This follows because the state \mathbf{e}_0 is dominated by any other state in S . In particular, $\mathbf{e}_0 \leq \mathbf{x}(t_2 - t_1, \mathbf{e}_0)$. Hence, by quasi-monotonicity, we have $\mathbf{x}(t_1, \mathbf{e}_0) \leq \mathbf{x}(t_1, \mathbf{x}(t_2 - t_1, \mathbf{e}_0)) = \mathbf{x}(t_2, \mathbf{e}_0)$.

Furthermore, Lemma 6 guarantees that if $\mathbf{x}(0) \in S$ then $\mathbf{x}(t) \in S$ for all $t \geq 0$. Hence, by adding (3) for all $i \geq k$ and using the fact that $\|\mathbf{x}(t)\|_1 < \infty$ gives

$$\begin{aligned} \dot{v}_k(t) &= \lambda \left[2\epsilon \left\{ x_{k-1}(t)x_{k-1-g}(t) - \sum_{i=k}^{k+g-1} x_i(t) \left(x_{i-1-g}(t) \right. \right. \right. \\ &\quad \left. \left. \left. - x_{i-g}(t) \right) \right\} + (1-2\epsilon)(x_{k-1}(t))^2 \right] - x_k(t). \end{aligned} \quad (27)$$

Since $x_{i-1-g}(t) \geq x_{i-g}(t)$ we have from the above

$$\begin{aligned} \dot{v}_k &\leq \lambda \left[2\epsilon x_{k-1} x_{k-1-g} + (1-2\epsilon)(x_{k-1})^2 \right] - x_k \\ &= \lambda \left[\epsilon x_{k-1-g}^2 + (1-\epsilon)x_{k-1}^2 - \epsilon(x_{k-1-g} - x_{k-1})^2 \right] - x_k \\ &\leq \lambda \left[\epsilon x_{k-1-g}^2 + (1-\epsilon)x_{k-1}^2 \right] - x_k \\ &\leq \lambda x_{k-1-g}^2 - x_k, \end{aligned} \quad (28)$$

where the last inequality follows by using $x_{k-1-g} \geq x_{k-1}$.

Existence of the Fixed Point \mathbf{x}^* : To prove the existence of the fixed point \mathbf{x}^* satisfying (5) and (6), we first show that $\|x_i(t, \mathbf{e}_0)\|_1$ remains uniformly bounded for all $t \geq 0$. Note that this is a stronger result than $\|x_i(t, \mathbf{e}_0)\|_1 < \infty$ for each $t \geq 0$ which has already been established in Lemma 6.

Proposition 7. *For $g \geq 0, \lambda < 1$, let $\mathbf{z}^* \in S$ be defined as $z_i^* = 1$ for all $i \leq 0$ and $z_i^* = \lambda(z_{i-1-g}^*)^2$ for all $i \geq 1$. Then, we have $\mathbf{x}(t, \mathbf{e}_0) \leq \mathbf{z}^*$ for all $t \geq 0$ which implies that*

$$\|\mathbf{x}(t, \mathbf{e}_0)\|_1 \leq \|\mathbf{z}^*\|_1 = (g+1) \sum_{i \geq 1} \lambda^{2i-1}, \quad \forall t \geq 0. \quad (29)$$

PROOF. Let $\mathbf{x}(0) = \mathbf{e}_0 \leq \mathbf{z}^*$. Then, from (26) we have $\dot{x}_k(t) \geq 0$ for all $k \geq 1$ and for all $t \geq 0$. This further implies that $\dot{v}_k(t) = \sum_{i \geq k} \dot{x}_i(t) \geq 0, \forall k \geq 1, \forall t \geq 0$. Since $\mathbf{x}(0) \leq \mathbf{z}^*$, for $\mathbf{x}(t) > \mathbf{z}^*$ for some t there must exist $t_* < t$ such that $x_l(t_*) = z_l^*$ and $\dot{x}_l(t_*) > 0$ for some $l \geq 1$. Let m be the smallest component where the above two conditions are satisfied. Then, substituting $k = m$ in (28), we obtain

$$\begin{aligned} \dot{v}_m(t_*) &\leq \lambda x_{m-1-g}^2(t_*) - x_m(t_*) = \lambda x_{m-1-g}^2(t_*) - z_m^* \\ &= \lambda(x_{m-1-g}^2(t_*) - (z_{m-1-g}^*)^2) \leq 0, \end{aligned}$$

where the last inequality follows from the definition of m . Since we already know that $\dot{v}_m(t_*) \geq 0$, the above inequality implies

$\dot{v}_m(t_*) = 0$. Hence, $\dot{x}_m(t_*) = \dot{v}_m(t_*) - \dot{v}_{m+1}(t_*) = -\dot{v}_{m+1}(t_*) \leq 0$ which contradicts the fact that $\dot{x}_m(t_*) > 0$. \square

Lemma 8. *Given $g \geq 0$, $\lambda < 1$, there exists $\mathbf{x}^* \in S$ such that*

$$\|\mathbf{x}(t, \mathbf{e}_0) - \mathbf{x}^*\|_1 \rightarrow 0, t \rightarrow \infty, \quad (30)$$

and $\mathbf{G}(\mathbf{x}^*) = \mathbf{0}$. Furthermore, \mathbf{x}^* satisfies (5) and (6).

PROOF. Since $x_i(t, \mathbf{e}_0) \in [0, 1]$ for each i and all $t \geq 0$ and $x_i(t, \mathbf{e}_0)$ is monotonically non-decreasing in time we must have $x_i(t) \rightarrow x_i^*$ as $t \rightarrow \infty$ for each $i \geq 1$ for some $\mathbf{x}^* = (x_i^*) \in \bar{S}$.

We first show that the component-wise limit \mathbf{x}^* defined above is also the ℓ_1 limit of $\mathbf{x}(t, \mathbf{e}_0)$ which will also imply that $\mathbf{x}^* \in S$. To show this, we note from Proposition 7 that the uniform bound on $\sum_{i \geq 1} x_i(t, \mathbf{e}_0)$ in (29) implies by dominated convergence theorem that

$$\lim_{t \rightarrow \infty} \|\mathbf{x}(t)\|_1 = \sum_{i \geq 1} \lim_{t \rightarrow \infty} x_i(t) = \sum_{i \geq 1} x_i^* = \|\mathbf{x}^*\|_1 \leq (g+1) \sum_{i \geq 1} \lambda^{2^i - 1}.$$

This shows that $\|\mathbf{x}(t, \mathbf{e}_0) - \mathbf{x}^*\|_1 \rightarrow 0$ as $t \rightarrow \infty$, and $\mathbf{x}^* \in S$.

It now remains to show that $\mathbf{G}(\mathbf{x}^*) = \mathbf{0}$. Note that the convergence of $\mathbf{x}(t) \rightarrow \mathbf{x}^*$ in ℓ_1 as $t \rightarrow \infty$, and the monotonicity of $\mathbf{x}(t)$ imply that for any $\delta > 0$ there exists a $t_\delta > 0$ such that for all $t \geq t_\delta$ we have

$$\begin{aligned} \delta \geq \|\mathbf{x}(t+h) - \mathbf{x}(t)\|_1 &\geq x_i(t+h) - x_i(t) = \int_t^{t+h} G_i(\mathbf{x}(s)) ds \\ &\geq h G_i(\mathbf{x}(t_h^*)), \forall i \geq 1, \forall h \geq 0, \end{aligned}$$

where $t_h^* \in [t, t+h]$ is the time at which the continuous function $G_i(\mathbf{x}(s))$ attains its minimum value in the compact interval $[t, t+h]$. Therefore, we have

$$G_i(\mathbf{x}(t_h^*)) \leq \frac{\delta}{h}, i \geq 1. \quad (31)$$

Now we can write

$$\begin{aligned} G_i(\mathbf{x}^*) &= G_i(\mathbf{x}^*) - G_i(\mathbf{x}(t_h^*)) + G_i(\mathbf{x}(t_h^*)) \\ &\leq \|G(\mathbf{x}^*) - G(\mathbf{x}(t_h^*))\|_1 + G_i(\mathbf{x}(t_h^*)) \leq L_\lambda^\epsilon \delta + \frac{\delta}{h}, \end{aligned} \quad (32)$$

where for the second inequality we use and the fact that the function G is Lipschitz with constant L_λ^ϵ and (31). Note that the above inequality is true for any $\delta > 0$. Therefore, by fixing $h > 0$ and letting $\delta \rightarrow 0$ we have $G_i(\mathbf{x}^*) = 0$ for all $i \geq 1$. Hence, $\mathbf{G}(\mathbf{x}^*) = \mathbf{0}$. Finally, we obtain (5) by using $\sum_{i \geq 1} G_i(\mathbf{x}^*) = 0$ and (6) by using $\sum_{i \geq k} G_i(\mathbf{x}^*) = 0$. \square

Global Stability and Uniqueness of the Fixed Point \mathbf{x}^* : Now we prove that for any $\mathbf{u} \in S$, $\mathbf{x}(t, \mathbf{u})$ converges to \mathbf{x}^* as $t \rightarrow \infty$ in ℓ_1 , where \mathbf{x}^* is the limit of $\mathbf{x}(t, \mathbf{e}_0)$ as defined in Lemma 8. By Proposition 7 and the dominated convergence theorem, it suffices to establish this convergence component-wise. Furthermore, it is sufficient to consider initial points $\mathbf{u} \leq \mathbf{x}^*$ and $\mathbf{u} \geq \mathbf{x}^*$ since, by the quasi-monotonicity of \mathbf{x} , we have $\mathbf{x}(t, \min(\mathbf{u}, \mathbf{x}^*)) \leq \mathbf{x}(t, \mathbf{u}) \leq \mathbf{x}(t, \max(\mathbf{u}, \mathbf{x}^*))$, where the min and the max are taken component-wise.

Consider the case when $\mathbf{x}(0) = \mathbf{u} \leq \mathbf{x}^*$. Since $\mathbf{u} \geq \mathbf{e}_0$, by the quasi-monotonicity of \mathbf{x} , we have $\mathbf{x}(t, \mathbf{e}_0) \leq \mathbf{x}(t, \mathbf{u}) \leq \mathbf{x}^*$, $\forall t \geq 0$. Hence, by Lemma 8, we have $\mathbf{x}(t, \mathbf{u}) \rightarrow \mathbf{x}^*$ since $\mathbf{x}(t, \mathbf{e}_0) \rightarrow \mathbf{x}^*$.

Next we consider the case where $\mathbf{x}(0) = \mathbf{u} \geq \mathbf{x}^*$. We first show that $v_k(t, \mathbf{u})$ remains uniformly bounded for all $t \geq 0$ and for all

$k \geq 1$. From quasi-monotonicity of \mathbf{x} it follows that $\mathbf{x}(t, \mathbf{u}) \geq \mathbf{x}^*$ for all $t \geq 0$. This implies, in particular, that $x_1(t, \mathbf{u}) \geq x_1^* = \lambda$. Hence, from (27) for $k = 1$, we have $\dot{v}_1(t, \mathbf{x}(0)) = \lambda - x_1(t, \mathbf{u}) \leq 0$. This implies that $0 \leq v_k(t, \mathbf{u}) \leq v_1(t, \mathbf{u}) \leq v_1(\mathbf{u})$ for all $t \geq 0$ and all $k \geq 1$.

We shall now establish the convergence $x_i(t, \mathbf{u}) \rightarrow x_i^*$ for all $i \geq 1$ by showing

$$\int_0^\infty (x_i(t, \mathbf{u}) - x_i^*) dt < C_i, \quad (33)$$

where $C_i > 0$ is a finite constant for each $i \geq 1$. To prove (33), we use induction on i . For $i = 1$, using (5) we have

$$\int_0^\tau (x_1(t, \mathbf{u}) - x_1^*) dt = \int_0^\tau (x_1(t, \mathbf{u}) - \lambda) dt = v_1(\mathbf{u}) - v_1(\tau, \mathbf{u}) \leq v_1(\mathbf{u}),$$

where the second equality follows from (27) for $k = 1$ and the inequality follows as $v_1(t, \mathbf{u})$ is uniformly bounded in t . Since the RHS is independent of τ , the integral on the left hand side must be bounded by $v_1(\mathbf{u})$ as $\tau \rightarrow \infty$. This shows the base case of the induction.

Now assume that (33) is true for all $i \leq L-1$. For $i = L$, using (27) and (6), we have

$$\begin{aligned} \int_0^\tau (x_L(t) - x_L^*) dt &\leq v_L(\mathbf{u}) + \lambda(1-2\epsilon) \int_0^\tau (x_{L-1}^2(t) - (x_{L-1}^*)^2) dt \\ &\quad + \lambda 2\epsilon \int_0^\tau (x_{L-1}(t)x_{L-1-g}(t) - x_{L-1}^*x_{L-1-g}^*) dt - 2\epsilon\lambda \\ &\times \int_0^\tau \sum_{i=L}^{L+g-1} \left[x_i(t)(x_{i-1-g}(t) - x_{i-g}(t)) - x_i^*(x_{i-1-g}^* - x_{i-g}^*) \right] dt, \end{aligned}$$

where the inequality follows from the uniform boundedness of $v_1(t, \mathbf{u})$ in t . To complete the proof, we shall now bound each integral term appearing on the RHS. By using the induction hypothesis and the inequalities $a^2 - b^2 \leq 2(a-b)$ and $ab - cd \leq (a-c) + (b-d)$ for $1 \geq a \geq c \geq 0$ and $1 \geq b \geq d \geq 0$, the integrals in the second and the third terms on the RHS can be easily bounded by $2C_{L-1}$ and $C_{L-1} + C_{L-1-g}$, respectively.

It now remains to bound the integral in the last term. Note that the third term contains $x_i(t)$ for $i \in \{L, L+1, \dots, L+g-1\}$ for which the induction hypothesis does not apply. Hence, to bound the integral we need to bound these terms. We note that by monotonicity we have $x_i(t) \geq x_i^*$ for all $i \geq 1$. Hence, $x_i(t)(x_{i-1-g}(t) - x_{i-g}(t)) - x_i^*(x_{i-1-g}^* - x_{i-g}^*) \geq x_i^*[(x_{i-1-g}(t) - x_{i-1-g}^*) - (x_{i-g}(t) - x_{i-g}^*)]$. Hence, the last term can be bounded above by

$$\begin{aligned} -\lambda 2\epsilon \sum_{i=L}^{L+g-1} x_i^* \left[\int_0^\tau (x_{i-1-g}(t) - x_{i-1-g}^*) dt \right. \\ \left. - \int_0^\tau (x_{i-g}(t) - x_{i-g}^*) dt \right]. \end{aligned}$$

Using the induction hypothesis, we can further bound this by $2\epsilon\lambda \sum_{i=L}^{L+g-1} (C_{i-1-g} + C_{i-g})$. This completes the proof of global stability of \mathbf{x}^* . Since all the trajectories converge to \mathbf{x}^* , it must be the unique solution of $\mathbf{G}(\mathbf{s}) = \mathbf{0}$ since starting from any other $\mathbf{y} \neq \mathbf{x}^*$, satisfying $\mathbf{G}(\mathbf{y}) = \mathbf{0}$, the trajectory remains at \mathbf{y} which contradicts the global stability of \mathbf{x}^* .

Limit Interchange: Note that Theorem 1.(i), implies that $\pi_n(S) = 1, \forall n$. Therefore, we have $\pi_n(\bar{S}) = 1$ for all n . Since, the space \bar{S} is compact, by Prohorov's theorem the sequence $(\pi_n)_n$ must converge weakly to the limit π^* with $\pi^*(\bar{S}) = 1$. Furthermore, since by (1), $\mathbb{E}_{\pi_n} \left[\sum_{i \geq 1} x_i^{(n)}(\infty) \right]$ is uniformly bounded in n , we have $\pi^*(S) = 1$. Now we prove that the measure π^* is the stationary measure of the mean-field process \mathbf{x} defined in (3). We know that $(\pi_n)_n \Rightarrow \pi^*$ and the space S is separable. Therefore, the Skorokhod's Representation Theorem implies that $\mathbf{x}^{(n)}(0) \xrightarrow{a.s.} \mathbf{x}(0)$. Moreover, if we start the process $\mathbf{x}^{(n)}(0) \sim \pi_n$, then $\mathbf{x}^{(n)}(t) \sim \pi_n$ for all $t \geq 0$. Hence, from Theorem 2.(i) it follows that $\mathbf{x}(t) \sim \pi^*$ for all $t \geq 0$. This proves that π^* is indeed the stationary measure for the mean-field process \mathbf{x} . Now from the global stability of the fixed point \mathbf{x}^* , it follows immediately that the stationary measure π^* is unique and is equal to $\delta_{\mathbf{x}^*}$. This completes the proof of limit interchange.

5.3 Heavy Traffic Limit for the Po2-(g, ϵ) scheme

In last we prove Theorem 2.(iii). Note that from Lemma 8, we know $\mathbf{x}(t, \mathbf{e}_0)$ converges to \mathbf{x}^* as $t \rightarrow \infty$ in ℓ_1 and from Proposition 7 we have $\|\mathbf{x}(t, \mathbf{e}_0)\|_1 \leq (g+1) \sum_{i \geq 1} \lambda^{2i-1}$ for all $t \geq 0$. Therefore, we can write $\|\mathbf{x}^*\|_1 \leq (g+1) \sum_{i \geq 1} \lambda^{2i-1}$. Hence, dividing the above inequality both side with $-\log(1-\lambda)$, taking limit $\lambda \rightarrow 1^-$, we obtain the bound given in (7).

6 CONCLUSION AND FUTURE DIRECTIONS

In this paper, we analyzed the effects of load comparison errors on the performance of the Po2 scheme. We considered two models of error. For the load-dependent error model, we showed that the Po2 scheme retains its benefits over the random scheme in the heavy traffic limit $\lambda \rightarrow 1$ for all values of g and ϵ . For the load-independent error model, we have shown that the Po2 scheme retains its benefits over the random scheme only if the probability of error $\epsilon \leq 1/2$. We introduce a general framework using Lyapunov functions to prove stability of our schemes. We also use a new approach to establish the mean-field limit results as the fixed point does not admit a recursive solution.

There are many interesting directions for further research. We have analyzed the performance of the Po2-(g, ϵ) scheme assuming g to be constant independent of n . It will be interesting to see the effect of varying g as a function of n . Another direction is to study the effects of delay in receiving the queue length information at the dispatcher. A more explicit delay dependent error model can be considered. Here, the challenge will be to analyze the effect of the delay on the performance of the Po2 scheme.

REFERENCES

- [1] R. R. Weber, "On the optimal assignment of customers to parallel servers," *Journal of Applied Probability*, vol. 15, no. 2, pp. 406–413, 1978.
- [2] W. Winston, "Optimality of the shortest line discipline," *Journal of applied probability*, vol. 14, no. 1, pp. 181–189, 1977.
- [3] S. Bhambay and A. Mukhopadhyay, "Asymptotic optimality of speed-aware jsq for heterogeneous service systems," *Performance Evaluation*, vol. 157, p. 102320, 2022.
- [4] —, "Optimal load balancing in heterogeneous server systems," *20th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt)*, pp. 113–120, 2022.
- [5] G. Goren, S. Vargaftik, and Y. Moses, "Distributed dispatching in the parallel server model," *IEEE/ACM Transactions on Networking*, 2022.
- [6] —, "Stochastic coordination in heterogeneous load balancing systems," in *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, 2021, pp. 403–414.
- [7] Y. Lu, Q. Xie, G. Kliot, A. Geller, J. R. Larus, and A. Greenberg, "Join-idle-queue: A novel load balancing algorithm for dynamically scalable web services," *Perform. Eval.*, vol. 68, no. 11, p. 1056–1071, nov 2011.
- [8] A. L. Stolyar, "Pull-based load distribution among heterogeneous parallel servers: the case of multiple routers," *Queueing Systems*, vol. 85, no. 1-2, pp. 31–65, 2017.
- [9] M. Mitzenmacher, "The power of two choices in randomized load balancing," *PhD thesis, University of California at Berkeley*, 1996.
- [10] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich, "Queueing system with selection of the shortest of two queues: An asymptotic approach," *Problemy Peredachi Informatsii*, vol. 32, no. 1, pp. 20–34, 1996.
- [11] W. Tarreau, "Test driving power of two random choices load balancing," 2019. [Online]. Available: <https://www.haproxy.com/blog/power-of-two-load-balancing/>
- [12] O. Garrett, "Nginx and the 'power of two choices' load-balancing algorithm," 2018. [Online]. Available: <https://www.nginx.com/blog/>
- [13] M. Smith, "Rethinking netflix's edge load balancing," 2018. [Online]. Available: <https://netflixtechblog.com/netflix-edge-load-balancing-695308b5548c>
- [14] X. Zhou, N. Shroff, and A. Wierman, "Asymptotically optimal load balancing in large-scale heterogeneous systems with multiple dispatchers," *Performance Evaluation*, vol. 145, pp. Art–No, 2021.
- [15] Q. Burke, P. McDaniel, T. La Porta, M. Yu, and T. He, "Misreporting attacks against load balancers in software-defined networking," *Mobile networks and applications*, 2021.
- [16] M. Mitzenmacher, A. W. Richa, and R. Sitaraman, "The power of two random choices: A survey of techniques and results," *Handbook of randomized computing, Netherlands*, 255–312, 2001.
- [17] D. Los and T. Sauerwald, "Balanced allocations with the choice of noise," in *Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing*, 2022, pp. 164–175.
- [18] G. Nadiradze, "On achieving scalability through relaxation," *PhD thesis*, 2021.
- [19] M. Bramson, "Stability of join the shortest queue networks," *The Annals of Applied Probability*, vol. 21, no. 4, pp. 1568–1625, 2011.
- [20] S. Foss and N. Chernova, "On the stability of a partially accessible multi-station queue with state-dependent routing," *Queueing Systems*, vol. 29, pp. 55–73, 1998.
- [21] M. Bramson, Y. Lu, and B. Prabhakar, "Asymptotic independence of queues under randomized load balancing," *Queueing Systems*, vol. 71, no. 3, pp. 247–292, 2012.
- [22] —, "Randomized load balancing with general service time distributions," *ACM SIGMETRICS performance evaluation review*, vol. 38, no. 1, pp. 275–286, 2010.
- [23] H. Chen and H.-Q. Ye, "Asymptotic optimality of balanced routing," *Operations research*, vol. 60, no. 1, pp. 163–179, 2012.
- [24] S. T. Maguluri, R. Srikant, and L. Ying, "Heavy traffic optimal resource allocation algorithms for cloud computing clusters," *Performance Evaluation*, vol. 81, pp. 20–39, 2014.
- [25] D. Mukherjee, S. C. Borst, J. S. Van Leeuwen, and P. A. Whiting, "Asymptotic optimality of power-of-d load balancing in large-scale systems," *Mathematics of Operations Research*, vol. 45, no. 4, pp. 1535–1571, 2020.
- [26] A. Budhiraja, D. Mukherjee, and R. Wu, "Supermarket model on graphs," *The Annals of Applied Probability*, vol. 29, no. 3, pp. 1740–1777, 2019.
- [27] D. Rutten and D. Mukherjee, "Load balancing under strict compatibility constraints," ser. SIGMETRICS '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 51–52.
- [28] Z. Zhao, D. Mukherjee, and R. Wu, "Exploiting data locality to improve performance of heterogeneous server clusters," *arXiv preprint arXiv:2211.16416*, 2022.
- [29] A. Mukhopadhyay and R. R. Mazumdar, "Analysis of randomized join-the-shortest-queue (jsq) schemes in large heterogeneous processor-sharing systems," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 2, pp. 116–126, 2015.
- [30] A. Mukhopadhyay, A. Karthik, and R. R. Mazumdar, "Randomized assignment of jobs to servers in heterogeneous clusters of shared servers for low delay," *Stochastic Systems*, vol. 6, no. 1, pp. 90–131, 2016.
- [31] S. Bhambay, A. Mukhopadhyay, and T. Vasantam, "The power of two choices with load comparison errors," 2023. [Online]. Available: <https://arxiv.org/abs/2303.00869>
- [32] F. Kelly and E. Yudovina, *Stochastic Networks*. Cambridge University Press, 2014.
- [33] P. W. Glynn and A. Zeevi, "Bounding stationary expectations of markov processes," in *Markov processes and related topics: a Festschrift for Thomas G. Kurtz*. Institute of Mathematical Statistics, 2008, pp. 195–214.
- [34] R. Srikant and L. Ying, *Communication networks: an optimization, control, and stochastic networks perspective*. Cambridge University Press, 2013.
- [35] S. N. Ethier and T. G. Kurtz, *Markov processes: characterization and convergence*. John Wiley & Sons, 2009, vol. 282.
- [36] T. G. Kurtz, *Approximation of population processes*. SIAM, 1981.
- [37] K. Deimling, *Ordinary differential equations in Banach spaces*. Springer, 2006, vol. 596.