

Manuscript version: Author's Accepted Manuscript

The version presented in WRAP is the author's accepted manuscript and may differ from the published version or Version of Record.

Persistent WRAP URL:

<http://wrap.warwick.ac.uk/116192>

How to cite:

Please refer to published version for the most recent bibliographic citation information. If a published version is known of, the repository item page linked to above, will contain details on accessing it.

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions.

© 2015 Elsevier. Licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International <http://creativecommons.org/licenses/by-nc-nd/4.0/>.



Publisher's statement:

Please refer to the repository item page, publisher's statement section, for further information.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk.

Mean Field and Propagation of Chaos in Multi-Class Heterogeneous Loss Models

Arpan Mukhopadhyay, A. Karthik, Ravi R. Mazumdar, Fabrice Guillemin

Abstract

We consider a system consisting of N parallel servers, where jobs with different resource requirements arrive and are assigned to the servers for processing. Each server has a finite resource capacity and therefore can serve only a finite number of jobs at a time. We assume that different servers have different resource capacities. A job is accepted for processing only if the resource requested by the job is available at the server to which it is assigned. Otherwise, the job is discarded or blocked. We consider randomized schemes to assign jobs to servers with the aim of reducing the average blocking probability of jobs in the system. In particular, we consider a scheme that assigns an incoming job to the server having maximum available vacancy or unused resource among d randomly sampled servers. We consider the system in the limit where both the number of servers and the arrival rates of jobs are scaled by the same large factor. This gives rise to a mean field analysis. We show that in the limiting system the servers behave independently – a property termed as *propagation of chaos*. Stationary probabilities of server occupancies are obtained from the stationary solution of the mean field. We further characterize the rate of decay of the stationary tail probabilities. Numerical results suggest that the proposed scheme significantly reduces the average blocking probability of jobs as compared to static schemes that probabilistically route jobs to servers independently of their states.

I. INTRODUCTION

Consider jobs with different resource requirements, arriving at a multi-server system consisting of a large number of parallel servers. Each server has a finite resource capacity and therefore

A. Mukhopadhyay, and R. R. Mazumdar are with the department of Electrical and Computer Engineering, University of Waterloo, Canada. Email: {arpan.mukhopadhyay, mazum}@uwaterloo.ca

A. Karthik is with Qualcomm, Bangalore, India. Email: k4ananth@uwaterloo.ca

F. Guillemin is with Orange Labs, France. Email: fabrice.guillemin@orange.com

A preliminary version of this work was presented at the 27th International Teletraffic Congress (ITC 27), Belgium, 2015.

can process only a finite number of jobs at a time. Different servers are assumed to have different capacities. Upon arrival, each job is routed/assigned to a server where it is either accepted or blocked depending on the availability of resource requested by the job at the server. If accepted, the processing of the job begins immediately at the server. The objective is to design job routing/assigning schemes that reduce the average blocking probability of jobs in the system.

Such a model arises frequently in the context of cloud computing systems that provide infrastructure as a service [1], [2]. A cloud service provider sells computing resources to its users in terms of virtual machines (VM's), that are computing instances consisting of various resources such as CPU, memory, storage etc. To meet different user demands the cloud operator allows users to choose from various types of VM's which differ by the amounts of resources they hold. We model situations [3], [4], where there is one bottleneck resource (e.g. memory). Hence, in our case the different VM's correspond to different amounts of the same resource. Each user requests a VM of specific type (e.g. large or small) for a required amount of time. The VM request is then assigned to a physical machine (PM) or server where the request is either accepted or blocked depending on the availability of resource at the server. If accepted, the user holds the VM for the duration of its service after which it is released. Therefore, to maintain a certain quality of service, a cloud service provider should aim at reducing the average blocking probability of users which measures the fraction of time a user is denied its requested resources.

The problem can be cast as a stochastic knapsack or a bin packing problem, where the blocking behavior is determined by the policy that assigns jobs or VM requests to servers. In particular, we consider a model where each server has the capacity to serve a finite number VM requests at a time and different servers can have different capacities. We refer to this as the *heterogeneous loss system* model. There are other possible abstractions to model clouds that involve buffering of jobs in an infinite queue as was analyzed in [5], [6]. However the loss model is particularly relevant for the Infrastructure-as-a-service (IaaS) paradigm offered by Amazon's EC2 [1] and Microsoft's Azure [2] where each server can process only a finite number of jobs at a time.

We consider a randomized scheme to assign VM requests to servers based on random sampling of d servers from the entire system. We show that assigning each request to the server having the maximum vacancy amongst the $d \geq 2$ sampled servers yields dramatic reduction in the blocking probability as compared to that in static assignment schemes ($d = 1$) where assignments are made independent of server states. Although, comparing the states of all servers would ideally

result in the best performance, such comparison involve high communication overhead due to large size of server farms. Moreover, we show that sampling too many servers is unnecessary since sampling only a few servers results in nearly optimal in terms of minimizing the average blocking probability of requests.

Related Literature: The routing scheme that we consider is a loss model analog of the *power-of- d* scheme considered in [7]–[9] for first-come-first-serve (FCFS) queues and in [10] for heterogeneous processor sharing (PS) servers. In the *power-of- d* scheme, each incoming job is assigned to the server having the least number of unfinished jobs among d servers, sampled randomly at the arrival instant of the job. Turner [11], [12] studied this scheme for a system of Erlang servers having infinite capacities in the large system limit. It was shown that in the large system limit, the system behavior can be characterized by a *mean field limit*, which satisfies a system of differential equations. The resulting tail distribution of server occupancies was shown to have a fast rate of decay even for small values of d . However, the existence and uniqueness of the equilibrium point of the mean field were not established. Xie *et al.* [4] recently analyzed the the *power-of- d* scheme in an Erlang loss system with identical (homogeneous) servers in the mean field limit. Existence, uniqueness, and global asymptotic stability of the equilibrium point of the mean field were established when there is a single class of customers in the system. For the case with multiple classes of customers, the authors derived a recursive relationship among the tail probabilities of the number of occupied resource units at each server assuming asymptotic independence of the servers in the limiting system. In the current paper, we generalize these results to the scenario where the servers have heterogeneous capacities.

Contributions: In this paper, we analyze the performance of the *power-of- d* scheme for a multi-server system consisting of heterogeneous loss servers in the presence of multiple job classes using mean field techniques. The mean field limit, given by the solution of a system of ordinary differential equations, approximates the behavior of the system when the arrival rates of different classes of jobs and the number of servers in the system are scaled by the same large factor. We establish the existence of an equilibrium point of the mean field. We also establish that the equilibrium point is unique and globally asymptotically stable when there is a single class of jobs arriving into the system. For the multi-class case, we provide some case studies which strongly suggest that the uniqueness and global asymptotic stability of the equilibrium point may be true in general for the multi-class case. Furthermore, it is shown that in the limiting

system any finite set of servers behave independently of each other and if the equilibrium point is globally asymptotically stable then the stationary distribution of states of any given finite set of servers can be obtained by taking the product of distributions determined by the equilibrium point on the given set. To show the independence mentioned above in the heterogeneous case, we introduce the notion of *intra-type exchangeability* where exchangeability holds only among servers having the same capacity. A bound on the rate of decay of stationary tail distribution of server occupancies similar to that obtained in [4] is found in the heterogeneous case. Numerical results show that the power-of- d scheme significantly reduces the average blocking probability of jobs in the system as compared to the static routing schemes and is nearly optimal in terms of minimizing the average blocking probability.

The rest of the paper is organized as follows. In Section II, we introduce the system model and describe the routing scheme studied in this paper. We then present the main results in Section III. Section IV presents a detailed analysis of the randomized scheme. Section VI provides numerical results to compare different routing alternatives. Section VII concludes the paper with some remarks.

II. SYSTEM MODEL

We consider a system consisting of N parallel servers, where jobs or VM requests arrive and request necessary resource for processing. We assume that there is only one bottleneck resource (e.g. memory) since such situations occur often in practice [3]. Each server holds a finite amount of the resource. The servers are categorized into M different types based on their resource capacities. Let $\mathcal{J} = \{1, 2, \dots, M\}$ be the index set of server types. A server of type $j \in \mathcal{J}$ is assumed to hold C_j units of resource. Without loss of generality, we assume that the capacities are ordered as follows: $C_1 \leq C_2 \leq \dots \leq C_M$. Furthermore, the fraction of type- j servers in the system is assumed to be fixed and is denoted by $\gamma_j \in [0, 1]$ for all $j \in \mathcal{J}$. Clearly, we have $\sum_{j=1}^M \gamma_j = 1$.

Jobs or VM requests are categorized into L classes depending on their resource requirements. Class $l \in \mathcal{L} = \{1, 2, \dots, L\}$ VM requests require $A_l \geq 0$ units of resource and are assumed to arrive at the system according to a Poisson process with rate $N\lambda_l$ independent of the other classes. We denote by $\underline{A} = (A_1, A_2, \dots, A_L)$ the L -dimensional vector of resource requirements. Upon arrival, a job is routed to one of the N servers according to the following routing scheme:

Power-of- d scheme: Upon arrival of each VM request, $d \geq 2$ *potential destination* servers are sampled uniformly at random from the set of N servers. The actual *destination* server for the arriving request is then chosen to be the server having the maximum vacancy or the maximum units of unused resource among the sampled servers. Ties among (sampled) servers of the same type are broken uniformly at random and ties across server types are broken by selecting the server type with the highest index (highest capacity). For example, if there are two type- j servers and one type $i < j$ server having the maximum vacancy among the sampled set of d servers, then any one of the two type j servers is chosen to be the destination server with probability $1/2$.

The destination server accepts the VM request if the necessary resource is available to process the request. If accepted, processing of the request begins immediately. Otherwise, the request is discarded or blocked and lost. We say that a server is in state $\underline{n} = (n_1, n_2, \dots, n_L)$ when, for each $l \in \mathcal{L}$, there are n_l jobs of class l in progress at the server. Clearly, the set of admissible states for a type $j \in \mathcal{J}$ server is given by $\mathcal{S}_j = \{\underline{n} \in \mathbb{Z}_+^L : \underline{n} \cdot \underline{A} \leq C_j\}$, where \mathbb{Z}_+ denotes the set of all non-negative integers and $\underline{n} \cdot \underline{A} \triangleq \sum_{l=1}^L n_l A_l$. We define the set of blocking states $\mathcal{B}_j^{(l)}$ for class $l \in \mathcal{L}$ jobs at a server of type $j \in \mathcal{J}$ as the set of states in \mathcal{S}_j for which the vacancy or the number of unused resource units is less than A_l , i.e., $\mathcal{B}_j^{(l)} = \{\underline{n} \in \mathbb{Z}_+^L : C_j - A_l < \underline{n} \cdot \underline{A} \leq C_j\}$. Clearly, a class- l VM request is blocked only when all the d potential destination servers are in the blocking states for the arriving request. The service times of accepted job requests are assumed to be independent and exponentially distributed random variables with mean 1. The service times of jobs are assumed to be also independent of the inter-arrival times of the jobs. The resource held by a request is released immediately upon the completion of its service.

III. MAIN RESULTS

In this section, we state (without proof) the main results of this paper. Our results are asymptotic in the sense that they are derived in the limit as the system size $N \rightarrow \infty$ keeping the proportions $\gamma_j, j \in \mathcal{J}$, fixed. Such results are especially useful in the context of cloud computing systems since they typically run tens of thousands of servers. Without loss of much generality we assume that C_j and A_l are non-negative integers for each $j \in \mathcal{J}$ and $l \in \mathcal{L}$.

Main results: *For the model described in Section II, let the stationary probability that a server of type $j \in \mathcal{J}$ has at least k units of occupied resource be $P_{k,j}^{(N)}$. Then $P_{k,j}^{(N)}$ converges to $P_{k,j}$*

in the limit as $N \rightarrow \infty$, where $P_{k,j}$ is the solution of the following recursive relation:

$$k(P_{k,j} - P_{k+1,j}) = \sum_{l=1}^L \frac{A_l \lambda_l}{\gamma_j} \left[\left(\sum_{i=1}^j \gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j+1}^M \gamma_i P_{k-A_l+C_i-C_{j+1},i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j}^M \gamma_i P_{k-A_l+C_i-C_{j+1},i} \right)^d \right], \quad (1)$$

for $1 \leq k \leq C_j$, with $P_{k,j} = 1$ for $k \leq 0$, and $P_{C_{j+1},j} = 0$ for all $j \in \mathcal{J}$. Furthermore, in the limit as $N \rightarrow \infty$ the servers become mutually independent and their stationary occupancy distributions are insensitive to the service time distribution.

Remark 1 (Propagation of chaos): We note that for finite system size N , the states of the servers are not independent of each other since at every arrival instant states of some randomly sampled servers are compared. However, in the limiting system ($N \rightarrow \infty$) the servers become mutually independent. This is known as the *propagation of chaos* or *asymptotic independence* property.

Remark 2 (Blocking probability): Using the independence of servers stated above and the probabilities $P_{k,j}$ found by solving (1) we can compute the blocking probability $P_{\text{blocking}}^{(l)}$ of class- l requests in the limiting system as follows: The stationary probability that a server of type j is in the blocking state for a class l job is $P_{C_j-A_l+1,j}$ and the probability that it is sampled at an arrival instant is γ_j . Thus the total probability that a randomly sampled server is in the blocking state for class- l requests is $\sum_{j \in \mathcal{J}} \gamma_j P_{C_j-A_l+1,j}$. Since the servers in the limiting system are mutually independent, the probability that the class- l arrival is blocked is given by $P_{\text{blocking}}^{(l)} = \left(\sum_{j \in \mathcal{J}} \gamma_j P_{C_j-A_l+1,j} \right)^d$.

A lower bound on blocking probability: The blocking probability of requests averaged over all classes obtained from any given routing scheme can be lower bounded as follows. For an arbitrary job routing scheme, the blocking probability averaged over all classes is given by $P_{\text{blocking}}^{\text{avg}} = \frac{\sum_{l \in \mathcal{L}} \lambda_l P_{\text{blocking}}^{(l)}}{\sum_{l \in \mathcal{L}} \lambda_l}$, where $P_{\text{blocking}}^{(l)}$ denotes the blocking probability of class- l jobs under that scheme. By Little's law the average number customer in the system is given by $(1 - P_{\text{blocking}}^{\text{avg}})N \sum_{l \in \mathcal{L}} \lambda_l$. Now if type j servers can accommodate a maximum of $B_j = \max_{\underline{n} \in \mathcal{S}_j} \left(\sum_{l \in \mathcal{L}} n_l \right)$ jobs of all classes combined, then the average number of jobs in the entire system is upper bounded by $N \sum_{j \in \mathcal{J}} \gamma_j B_j$. We therefore have the following lower bound on the average blocking probability:

$$P_{\text{blocking}}^{\text{avg}} \geq \left(1 - \frac{\sum_{j \in \mathcal{J}} \gamma_j B_j}{\sum_{l \in \mathcal{L}} \lambda_l}\right)_+ = \left(1 - \frac{\lambda_{\text{crit}}}{\lambda}\right)_+, \quad (2)$$

where $\lambda_{\text{crit}} = \sum_{j \in \mathcal{J}} \gamma_j B_j$ is the critical load on the system, $\lambda = \sum_{l \in \mathcal{L}} \lambda_l$ is the total arrival rate per server, and $(w)_+ = \max(0, w)$. In Section VI, we compare the blocking probability for the power-of- d scheme with the lower bound derived above. We conclude that the power-of- d scheme is nearly optimal in terms of minimizing blocking probability even for small d .

If we specialize (1) to the case where only a single class of jobs ($L = 1$) requiring one unit of resource from all servers ($A = 1$) arrive according to a Poisson process with rate $N\lambda$, then (1) simplifies to the following recursive relation:

$$P_{k,j} - P_{k+1,j} = \frac{\lambda}{\gamma_j k} \left[\left(\sum_{i=1}^j \gamma_i P_{k-1+C_i-C_j,i} + \sum_{i=j+1}^M \gamma_i P_{k+C_i-C_j,i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-1+C_i-C_j,i} + \sum_{i=j}^M \gamma_i P_{k+C_i-C_j,i} \right)^d \right]. \quad (3)$$

In this above, $P_{k,j}$ is the stationary probability that there are at least k jobs in progress at a type j server in the limiting system. Using (3), explicit upper bounds on the rate of decay of the tail probabilities $P_{k,j}$, $k \in \{0, 1, 2, \dots, C_j\}$, $j \in \mathcal{J}$ can be obtained. This is done in the following proposition whose proof is similar to the proof of Theorem 2 of [4] for the homogeneous loss model.

Proposition 1: Let $\{\bar{P}_k, 0 \leq k \leq C_M\}$ be defined as follows: $\bar{P}_k = 1$ for $0 \leq k \leq k_0$ and

$$\bar{P}_k = \frac{\lambda^{d^{k-k_0}-1}}{([\lambda] + k - k_0)([\lambda] + k - k_0 - 1)^d \dots ([\lambda] + 1)^{d^{k-k_0-1}}}, \quad (4)$$

for $k_0 + 1 \leq k \leq C_M$, where $k_0 = \lfloor \lambda \rfloor + C_M - C_1$, and $\lfloor y \rfloor$ denotes the greatest integer not exceeding y . Then for the single class case where each job requires unit resource we have $\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \leq \bar{P}_k$ for $0 \leq k \leq C_M$. In particular, the average user blocking probability $P_{\text{blocking}}^{\text{avg}} = \left(\sum_{j \in \mathcal{J}} \gamma_j P_{C_j,j}\right)^d \leq \bar{P}_{C_M}^d$.

Proof: The proof is given in Appendix A. ■

The above proposition shows that for $d \geq 2$ the quantity $\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j}$ decreases with k at a rate much faster than that for $d = 1$. This shows the efficacy of sampling a small number

servers from the system over randomly sampling the destination server independent of the server states.

IV. MEAN FIELD ANALYSIS

In this section we provide detailed proofs of the main results discussed in Section III. An exact characterization of the stationary regime is difficult for finite N due to the fact that the arrival rate at a given server depends on the states of other servers. However, it is possible to analyze the system in the limit as the system size $N \rightarrow \infty$. Such a limit is known as the mean field limit [7], [15], [16]. We first introduce the notation and mathematical framework for the analysis.

Notations: The unit vector in \mathbb{Z}_+^L with one in the r^{th} position is denoted by \underline{e}_r and by $\underline{e} \in \mathbb{Z}_+^L$ we denote the L -dimensional vector of all ones. Further, for each $j \in \mathcal{J}$ we denote the space of probability distributions on \mathcal{S}_j by:

$$\mathcal{U}_j = \left\{ (g_{\underline{n}})_{\underline{n} \in \mathcal{S}_j} : g_{\underline{n}} \geq 0 \text{ for all } \underline{n} \in \mathcal{S}_j, \sum_{\underline{n} \in \mathcal{S}_j} g_{\underline{n}} = 1 \right\}.$$

The set of empirical probability distributions on \mathcal{S}_j when the system size is N is denoted by $\mathcal{U}_j^{(N)}$, i.e., $\mathcal{U}_j^{(N)} = \left\{ (g_{\underline{n}})_{\underline{n} \in \mathcal{S}_j} \in \mathcal{U}_j : N\gamma_j g_{\underline{n}} \in \mathbb{Z}_+ \right\}$. We will mainly be interested in the spaces $\mathcal{U} = \prod_{j \in \mathcal{J}} \mathcal{U}_j$ and $\mathcal{U}^{(N)} = \prod_{j \in \mathcal{J}} \mathcal{U}_j^{(N)}$, which are the Cartesian products of the spaces \mathcal{U}_j and $\mathcal{U}_j^{(N)}$, respectively, over $j \in \mathcal{J}$. A point in the space \mathcal{U} (or in $\mathcal{U}^{(N)}$) is denoted by $\mathbf{u} = (u_{\underline{n},j}, \underline{n} \in \mathcal{S}_j, j \in \mathcal{J})$ with the understanding that for each $j \in \mathcal{J}$ the collection $(u_{\underline{n},j}, \underline{n} \in \mathcal{S}_j)$ belongs to \mathcal{U}_j ($\mathcal{U}_j^{(N)}$). Since for each $j \in \mathcal{J}$, the set \mathcal{S}_j is finite and \mathcal{U}_j denotes the space of probability measures on \mathcal{S}_j , the space $\mathcal{U} = \prod_{j \in \mathcal{J}} \mathcal{U}_j$ is convex and compact under any norm and all norms defined on this space are equivalent. We shall be using both the Euclidean norm and the norm induced by the metric ρ (defined below) according to our convenience: for any two points \mathbf{u}, \mathbf{w} in \mathcal{U} define $\rho(\mathbf{u}, \mathbf{w}) = \sup_{j \in \mathcal{J}} \sup_{\underline{n} \in \mathcal{S}_j} \frac{|u_{\underline{n},j} - w_{\underline{n},j}|}{(\underline{n} \cdot \underline{e}) + 1}$.

For a measure space (H, \mathcal{H}, μ) and a μ -integrable function $f : H \rightarrow \mathbb{R}$, we define duality brackets as $\langle f, \mu \rangle = \int f d\mu$. Law of a random variable X is denoted by $\mathcal{L}(X)$. Weak convergence (convergence in distribution) of a sequence of probability measures ν_n (random variables X_n) to a probability measure ν (random variable X) is denoted by $\nu_n \Rightarrow \nu$ ($X_n \Rightarrow X$).

Analysis: For each $t \geq 0$ and $\underline{n} \in \mathcal{S}_j$, let $x_{\underline{n},j}^{(N)}(t)$ denote the fraction of type- j servers in state \underline{n} at time t . We define $\mathbf{x}^{(N)}(t) = \left(x_{\underline{n},j}^{(N)}(t), \underline{n} \in \mathcal{S}_j, j \in \mathcal{J} \right)$. Clearly, $(\mathbf{x}^{(N)}(t), t \geq 0)$ is a Markov

process with state space $\mathcal{U}^{(N)}$, i.e., for each $j \in \mathcal{J}$ the collection $\left(x_{\underline{n},j}^{(N)}(t), n \in \mathcal{S}_j\right)$ denotes the empirical probability distribution of states of type- j servers at time $t \geq 0$. The generator $\mathbf{A}^{(N)}$ of the Markov process $(\mathbf{x}^{(N)}(t), t \geq 0)$ acting on functions $\varphi : \mathcal{U}^{(N)} \rightarrow \mathbb{R}$ is given by $\mathbf{A}^{(N)}\varphi(\mathbf{u}) = \sum_{\mathbf{h} \neq \mathbf{u}} r(\mathbf{u} \rightarrow \mathbf{v}) (\varphi(\mathbf{v}) - \varphi(\mathbf{u}))$, where $r(\mathbf{u} \rightarrow \mathbf{v})$ denotes the transition rate from state $\mathbf{u} \in \mathcal{U}^{(N)}$ to state $\mathbf{v} \in \mathcal{U}^{(N)}$. In the following lemma, we provide the expression for the generator.

Lemma 1: *Let $\mathbf{u} \in \mathcal{U}^{(N)}$ and $\mathbf{e}(\underline{n}, j) = (e_{\underline{k},i}, \underline{k} \in \mathcal{S}_i, i \in \mathcal{J})$ be the unit vector with $e_{\underline{n},j} = 1$ and $e_{\underline{k},i} = 0$ if $\underline{k} \neq \underline{n}$ or $i \neq j$. Then the generator $\mathbf{A}^{(N)}$ of the process $\mathbf{x}^{(N)}(\cdot)$ acting on functions $\varphi : \mathcal{U}^{(N)} \rightarrow \mathbb{R}$ is given by*

$$\begin{aligned} \mathbf{A}^{(N)}\varphi(\mathbf{u}) = N \sum_{j \in \mathcal{J}} \sum_{\underline{n} \in \mathcal{S}_j} \sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \mathbf{u})}{E(\underline{n} - \underline{e}_l, j, \mathbf{u})} \gamma_j u_{\underline{n} - \underline{e}_l, j} \left(\varphi\left(\mathbf{u} - \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} + \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}\right) \right. \right. \\ \left. \left. - \varphi(\mathbf{u}) \right) + \gamma_j u_{\underline{n}, j} n_l \left(\varphi\left(\mathbf{u} + \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} - \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}\right) - \varphi(\mathbf{u}) \right) \right] I_{\mathcal{S}_j}(\underline{n} - \underline{e}_l), \quad (5) \end{aligned}$$

where $I_A(\cdot)$ denotes the indicator function on the set A and for $\underline{n} \in \mathcal{S}_j$ and $i \in \mathcal{J}$ we define

$$E(\underline{n}, i, j, \mathbf{u}) = \gamma_i \sum_{\substack{\underline{n}' \in \mathcal{S}_i: \\ \underline{n}' \cdot \underline{A} = \underline{n} \cdot \underline{A} + C_i - C_j}} u_{\underline{n}', i} \quad (6)$$

$$G(\underline{n}, i, j, \mathbf{u}) = \gamma_i \sum_{\substack{\underline{n}' \in \mathcal{S}_i: \\ \underline{n}' \cdot \underline{A} > \underline{n} \cdot \underline{A} + C_i - C_j}} u_{\underline{n}', i} \quad (7)$$

$$GE(\underline{n}, i, j, \mathbf{u}) = G(\underline{n}, i, j, \mathbf{u}) + E(\underline{n}, i, j, \mathbf{u}). \quad (8)$$

and

$$\begin{aligned} F(\underline{n}, j, \mathbf{u}) = \left(\sum_{i=1}^j GE(\underline{n}, i, j, \mathbf{u}) + \sum_{i=j+1}^M G(\underline{n}, i, j, \mathbf{u}) \right)^d \\ - \left(\sum_{i=1}^{j-1} GE(\underline{n}, i, j, \mathbf{u}) + \sum_{i=j}^M G(\underline{n}, i, j, \mathbf{u}) \right)^d. \quad (9) \end{aligned}$$

Proof: The proof is given in Appendix B. ■

Using the generator $\mathbf{A}^{(N)}$, we now show that as $N \rightarrow \infty$, the sequence of processes $\{(\mathbf{x}^{(N)}(t), t \geq 0)\}_N$ converges to a deterministic process known as the *mean field limit* of the system.

Theorem 1: *If $\mathbf{x}^{(N)}(0)$ converges in distribution to some constant $\mathbf{u}_0 \in \mathcal{U}$ as $N \rightarrow \infty$, then as $N \rightarrow \infty$ the process $(\mathbf{x}^{(N)}(t), t \geq 0)$ converges in distribution to a deterministic process $(\mathbf{x}(t), t \geq 0)$, which takes values in the space \mathcal{U} and is given by the unique solution of the following system of differential equations*

$$\mathbf{x}(0) = \mathbf{u}_0, \quad (10)$$

$$\dot{\mathbf{x}}(t) = \mathbf{h}(\mathbf{x}(t)), \quad (11)$$

where $\mathbf{h}(\cdot) = (h_{\underline{n},j}(\cdot), j \in \mathcal{J}, \underline{n} \in \mathcal{S}_j)$ and for each $\underline{n} \in \mathcal{S}_j$ and $j \in \mathcal{J}$

$$h_{\underline{n},j}(\mathbf{x}) = \sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \mathbf{x})}{E(\underline{n} - \underline{e}_l, j, \mathbf{x})} x_{\underline{n} - \underline{e}_l, j} - n_l x_{\underline{n}, j} \right] I_{\mathcal{S}_j}(\underline{n} - \underline{e}_l) - \left[\lambda_l \frac{F(\underline{n}, j, \mathbf{x})}{E(\underline{n}, j, \mathbf{x})} x_{\underline{n}, j} - (n_l + 1) x_{\underline{n} + \underline{e}_l, j} \right] I_{\mathcal{S}_j}(\underline{n} + \underline{e}_l). \quad (12)$$

In the above equation, $I_A(\cdot)$ denotes indicator function of the set A and $E(\underline{n}, i, j, \mathbf{u})$, $F(\underline{n}, j, \mathbf{u})$ are as defined in Lemma 1.

Proof: The proof is given in Appendix C ■

To emphasize the dependence of the mean field process $(\mathbf{x}(t), t \geq 0)$ on the initial value, we will sometimes denote the process with $\mathbf{x}(0) = \mathbf{u}_0$ as $(\mathbf{x}(t, \mathbf{u}_0), t \geq 0)$. It is important to characterize the properties of the points $\boldsymbol{\pi} = (\pi_{\underline{n},j}, \underline{n} \in \mathcal{S}_j, j \in \mathcal{J})$ satisfying $\mathbf{h}(\boldsymbol{\pi}) = \mathbf{0}$. Such points are called the *equilibrium points* of the mean field since $\mathbf{x}(t, \boldsymbol{\pi}) = \boldsymbol{\pi}$ for all $t \geq 0$ and describes the stationary behavior of the limiting system. Hence, by definition, a point $\boldsymbol{\pi} \in \mathcal{U}$ is an equilibrium point of the mean field $(\mathbf{x}(t), t \geq 0)$ if and only if it satisfies $h_{\underline{n},j}(\boldsymbol{\pi}) = 0$, i.e., (from (12))

$$\begin{aligned} \sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \boldsymbol{\pi})}{E(\underline{n} - \underline{e}_l, j, j, \boldsymbol{\pi})} \pi_{\underline{n} - \underline{e}_l, j} - n_l \pi_{\underline{n}, j} \right] I_{\mathcal{S}_j}(\underline{n} - \underline{e}_l) \\ = \sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n}, j, j, \boldsymbol{\pi})}{E(\underline{n}, j, j, \boldsymbol{\pi})} \pi_{\underline{n}, j} - (n_l + 1) \pi_{\underline{n} + \underline{e}_l, j} \right] I_{\mathcal{S}_j}(\underline{n} + \underline{e}_l), \end{aligned} \quad (13)$$

for each $j \in \mathcal{J}$ and $\underline{n} \in \mathcal{S}_j$. In the next theorem, we show that such a point indeed exists.

Theorem 2: There exists an equilibrium point $\boldsymbol{\pi}$ of the mean field process $(\mathbf{x}(t), t \geq 0)$ in the space \mathcal{U} .

Proof: Consider a point $\mathbf{x} \in \mathcal{U}$. For each $j \in \mathcal{J}$, $l \in \mathcal{L}$ and $\underline{n} \in \mathcal{S}_j$, define

$$\lambda_{\underline{n}, j}^{(l)}(\mathbf{x}) = \lambda_l \frac{F(\underline{n}, j, \mathbf{x})}{E(\underline{n}, j, j, \mathbf{x})} > 0. \quad (14)$$

Next, we define the quantities $y_{\underline{n}, j}(\mathbf{x})$, for $j \in \mathcal{J}$, $\underline{n} \in \mathcal{S}_j$ as the solution to the following system of linear equations

$$\begin{aligned} \sum_{l \in \mathcal{L}} \left[\lambda_{\underline{n} - \underline{e}_l, j}^{(l)}(\mathbf{x}) y_{\underline{n} - \underline{e}_l, j}(\mathbf{x}) - n_l y_{\underline{n}, j}(\mathbf{x}) \right] I_{\mathcal{S}_j}(\underline{n} - \underline{e}_l) \\ = \sum_{l \in \mathcal{L}} \left[\lambda_{\underline{n}, j}^{(l)}(\mathbf{x}) y_{\underline{n}, j}(\mathbf{x}) - (n_l + 1) y_{\underline{n} + \underline{e}_l, j}(\mathbf{x}) \right] I_{\mathcal{S}_j}(\underline{n} + \underline{e}_l), \end{aligned} \quad \text{for } j \in \mathcal{J} \text{ and } \underline{n} \in \mathcal{S}_j \quad (15)$$

and $\sum_{\underline{n} \in \mathcal{S}_j} y_{\underline{n}, j}(\mathbf{x}) = 1$ for each $j \in \mathcal{J}$. Clearly, the solution $\mathbf{y}(\mathbf{x}) = (y_{\underline{n}, j}(\mathbf{x}), \underline{n} \in \mathcal{S}_j, j \in \mathcal{J})$ to the above set of linear equations satisfies

$$\lambda_{\underline{n} - \underline{e}_l, j}^{(l)}(\mathbf{x}) y_{\underline{n} - \underline{e}_l, j}(\mathbf{x}) I_{\mathcal{S}_j}(\underline{n} - \underline{e}_l) = n_l y_{\underline{n}, j}(\mathbf{x}) \text{ for all } \underline{n} \in \mathcal{S}_j, j \in \mathcal{J}. \quad (16)$$

The above equations (which imply that $y_{\underline{n}, j}(\mathbf{x})$ has the same sign for each $\underline{n} \in \mathcal{S}_j$ and $j \in \mathcal{J}$) together with $\sum_{\underline{n} \in \mathcal{S}_j} y_{\underline{n}, j}(\mathbf{x}) = 1$ imply that $\mathbf{y}(\mathbf{x}) \in \mathcal{U}$ for all $\mathbf{x} \in \mathcal{U}$. Furthermore, it is easy to see that the map $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$, as defined above, is continuous on the space \mathcal{U} (since $(y_{\underline{n}, j}(\mathbf{x}), j \in \mathcal{J}, \underline{n} \in \mathcal{S}_j)$ is continuous in $(\lambda_{\underline{n}, j}^{(l)}(\mathbf{x}), j \in \mathcal{J}, \underline{n} \in \mathcal{S}_j, l \in \mathcal{L})$ which in turn is continuous in \mathbf{x}). Since \mathcal{U} is convex and compact, Brouwer's fixed point theorem guarantees the existence of a fixed point of the map $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$. From (15) and the definition of the quantities $(\lambda_{\underline{n}, j}^{(l)}(\mathbf{x}), j \in \mathcal{J}, \underline{n} \in \mathcal{S}_j, l \in \mathcal{L})$ as given by (14), it is clear that any fixed point $\boldsymbol{\pi}$ of the map

$\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$ satisfies (13) and hence is an equilibrium point of the mean field process. This proves the existence of an equilibrium point $\boldsymbol{\pi}$ in \mathcal{U} of the mean field $(\mathbf{x}(t), t \geq 0)$. ■

Remark 3: We note that for each $\mathbf{x} \in \mathcal{U}$, the mapping $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$ satisfies (16). Hence, the fixed point $\boldsymbol{\pi}$, which is also the equilibrium point of the mean field, satisfies

$$\lambda_{\underline{n}-\underline{e}_l, j}^{(l)}(\boldsymbol{\pi}) \pi_{\underline{n}-\underline{e}_l, j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = n_l \pi_{\underline{n}, j} \text{ for all } \underline{n} \in \mathcal{S}_j, j \in \mathcal{J}. \quad (17)$$

We shall use this fact later in Section V.

We now focus on the single class case ($L = 1$) and show that the equilibrium point of the mean field is *unique* and *globally asymptotically stable*, i.e., for any $\mathbf{x}(0) \in \mathcal{U}$ we have $\lim_{t \rightarrow \infty} \mathbf{x}(t) = \boldsymbol{\pi}$, where $\boldsymbol{\pi}$ denotes the equilibrium point of the mean field.

Theorem 3: For the single class case ($L = 1$), the mean field $(\mathbf{x}(t), t \geq 0)$ has a unique, globally asymptotically stable equilibrium point $\boldsymbol{\pi} \in \mathcal{U}$.

Proof: We note that the uniqueness of the equilibrium point follows (by the uniqueness of limit) if one can show that any equilibrium point $\boldsymbol{\pi}$ is globally asymptotically stable. Hence, it is sufficient to establish the global asymptotic stability of any equilibrium point of the mean field for the single class case.

For the single class case, we assume without loss of generality that all incoming jobs require one unit of resource and they arrive according to a Poisson process with rate $N\lambda$, i.e., $A_1 = 1$ and $\lambda_1 = \lambda$. Hence, $\mathcal{S}_j = \{0, 1, \dots, C_j\}$, $\mathcal{U}_j = \left\{ (g_n)_{n \in \mathcal{S}_j} : g_n \geq 0, \sum_{n \in \mathcal{S}_j} g_n = 1 \right\}$, $\mathcal{U} = \prod_{j \in \mathcal{J}} \mathcal{U}_j$. In this case, the mean field $(\mathbf{x}(t) = (x_{n,j}(t), n \in \mathcal{S}_j, j \in \mathcal{J}), t \geq 0)$ satisfies the following system of differential equations (from (10)-(11))

$$\mathbf{x}(0) = \mathbf{u}_0, \quad (18)$$

$$\dot{\mathbf{x}}(t) = \mathbf{h}(\mathbf{x}(t)), \quad (19)$$

where $\mathbf{h}(\cdot) = (h_{n,j}(\cdot), n \in \mathcal{S}_j, j \in \mathcal{J})$ is given by (specializing (12) to the case under consideration)

$$h_{n,j}(\mathbf{x}) = \left[\frac{\lambda}{\gamma_j} F(n-1, j, \mathbf{x}) - n x_{n,j} \right] I_{1 \leq n \leq C_j} - \left[\frac{\lambda}{\gamma_j} F(n, j, \mathbf{x}) - (n+1) x_{n+1,j} \right] I_{0 \leq n \leq C_j-1}. \quad (20)$$

The mean field can equivalently be expressed in terms of the tail sums $\tilde{x}_{k,j}(t) = \sum_{n=k}^{C_j} x_{n,j}(t)$, $k \in \mathcal{S}_j$, $j \in \mathcal{J}$. We define $\tilde{\mathbf{x}}(t) = (\tilde{x}_{k,j}(t), k \in \mathcal{S}_j, j \in \mathcal{J})$. Hence, from (19) and (20) we have

$$\tilde{\mathbf{x}}(0) = \tilde{\mathbf{u}}_0, \quad (21)$$

$$\dot{\tilde{\mathbf{x}}}(t) = \tilde{\mathbf{h}}(\tilde{\mathbf{x}}(t)), \quad (22)$$

where the mapping $\tilde{\mathbf{h}}(\cdot) = (\tilde{h}_{k,j}(\cdot), k \in \mathcal{S}_j, j \in \mathcal{J})$ is given by $\tilde{h}_{0,j}(\tilde{\mathbf{x}}) = 0$ for all $j \in \mathcal{J}$ and for $1 \leq k \leq C_j$

$$\begin{aligned} \tilde{h}_{k,j}(\tilde{\mathbf{x}}) = \frac{\lambda}{\gamma_j} & \left[\left(\sum_{i=1}^j \gamma_i \tilde{x}_{k-1+C_i-C_j,i} + \sum_{i=j+1}^M \gamma_i \tilde{x}_{k+C_i-C_j,i} \right)^d \right. \\ & \left. - \left(\sum_{i=1}^{j-1} \gamma_i \tilde{x}_{k-1+C_i-C_j,i} + \sum_{i=j}^M \gamma_i \tilde{x}_{k+C_i-C_j,i} \right)^d \right] - k (\tilde{x}_{k,j} - \tilde{x}_{k+1,j}) \end{aligned} \quad (23)$$

We say that $\tilde{\mathbf{u}} \leq \tilde{\mathbf{u}}'$ if $\tilde{u}_{k,j} \leq \tilde{u}'_{k,j}$ for all $k \in \mathcal{S}_j$ and $j \in \mathcal{J}$. We first prove the following monotonicity property of the mean field with respect to the initial condition.

Lemma 2: If $\tilde{\mathbf{u}}_0 \leq \tilde{\mathbf{u}}'_0$ then $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \leq \tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}'_0)$ for all $t \geq 0$.

Proof: Clearly, the right hand side of (23) is non-decreasing in $\tilde{x}_{n,i}(t)$ for all $(n, i) \neq (k, j)$. Hence, (23) defines a quasi-monotone system of differential equations. The proof of the lemma now follows directly from pages 70-74 of [17]. \blacksquare

We now define $z(t, \tilde{\mathbf{u}}_0) = \sum_{j \in \mathcal{J}} \gamma_j \sum_{k=1}^{C_j} \tilde{x}_{k,j}(t, \tilde{\mathbf{u}}_0)$. Clearly, $z(t, \tilde{\mathbf{u}}_0)$ denotes the mean number of customers in the limiting system at time t when the initial state is $\tilde{\mathbf{u}}_0$. Using (23) we obtain

$$\frac{dz(t, \tilde{\mathbf{u}}_0)}{dt} = \lambda \left(1 - \left(\sum_{j \in \mathcal{J}} \gamma_j \tilde{x}_{C_j,j}(t, \tilde{\mathbf{u}}_0) \right)^d \right) - z(t, \tilde{\mathbf{u}}_0). \quad (24)$$

Let $\tilde{\boldsymbol{\pi}}$ be an equilibrium point of the process $\tilde{\mathbf{x}}(\cdot)$. Hence, from (24) we have

$$\lambda \left(1 - \left(\sum_{j \in \mathcal{J}} \gamma_j \tilde{\pi}_{C_j,j} \right)^d \right) = z(t, \boldsymbol{\pi}) = \sum_{j \in \mathcal{J}} \gamma_j \sum_{k=1}^{C_j} \tilde{\pi}_{k,j} \quad (25)$$

Now, from Lemma 2 we have

$$\tilde{\mathbf{x}}(t, \min(\tilde{\mathbf{u}}_0, \tilde{\boldsymbol{\pi}})) \leq \tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \leq \tilde{\mathbf{x}}(t, \max(\tilde{\mathbf{u}}_0, \tilde{\boldsymbol{\pi}})), \quad (26)$$

where the maximum and the minimum are taken component-wise. Hence, to establish $\lim_{t \rightarrow \infty} \tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) = \tilde{\boldsymbol{\pi}}$ for all $\tilde{\mathbf{u}}_0$, it is sufficient to show that the convergence holds for $\tilde{\mathbf{u}}_0 \geq \tilde{\boldsymbol{\pi}}$ and for $\tilde{\mathbf{u}}_0 \leq \tilde{\boldsymbol{\pi}}$.

To show $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \rightarrow \tilde{\boldsymbol{\pi}}$ for $\tilde{\mathbf{u}}_0 \geq \tilde{\boldsymbol{\pi}}$ it is sufficient to show that

$$\int_0^\infty (\tilde{x}_{n,j}(t, \tilde{\mathbf{u}}_0) - \tilde{\pi}_{n,j}) dt < \infty, \text{ for all } j \in \mathcal{J}, 1 \leq n \leq C_j. \quad (27)$$

Similarly for $\tilde{\mathbf{u}}_0 \leq \tilde{\boldsymbol{\pi}}$ the convergence $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \rightarrow \tilde{\boldsymbol{\pi}}$ will follow if we can show that

$$\int_0^\infty (\tilde{\pi}_{n,j} - \tilde{x}_{n,j}(t, \tilde{\mathbf{u}}_0)) dt < \infty, \text{ for all } j \in \mathcal{J}, 1 \leq n \leq C_j \quad (28)$$

We discuss the proof for $\tilde{\mathbf{u}}_0 \geq \tilde{\boldsymbol{\pi}}$. The proof for $\tilde{\mathbf{u}}_0 \leq \tilde{\boldsymbol{\pi}}$ follows similarly.

For $\tilde{\mathbf{u}}_0 \geq \tilde{\boldsymbol{\pi}}$ we have using Lemma 2 that $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \geq \tilde{\boldsymbol{\pi}}$ for all $t \geq 0$. Hence, to prove (27) it is sufficient to show that $\int_0^\infty \left(z(t, \tilde{\mathbf{u}}_0) - \sum_{j \in \mathcal{J}} \gamma_j \sum_{n=1}^{C_j} \tilde{\pi}_{n,j} \right) dt < \infty$. We have

$$\begin{aligned} & \int_0^\tau \left(z(t, \tilde{\mathbf{u}}_0) - \sum_{j \in \mathcal{J}} \gamma_j \sum_{n=1}^{C_j} \tilde{\pi}_{n,j} \right) dt = - \int_0^\tau \frac{dz(t, \tilde{\mathbf{u}}_0)}{dt} dt \\ & - \int_0^\tau \lambda \left(\left(\sum_{j \in \mathcal{J}} \gamma_j \tilde{\pi}_{C_j,j} \right)^d - \left(\sum_{j \in \mathcal{J}} \gamma_j \tilde{x}_{C_j,j}(t, \tilde{\mathbf{u}}_0) \right)^d \right) dt \\ & \leq (z(0, \tilde{\mathbf{u}}_0) - z(\tau, \tilde{\mathbf{u}}_0)) \leq z(0, \tilde{\mathbf{u}}_0) \end{aligned}$$

where the first equality follows from (24) and (25); the second inequality follows since $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \geq \tilde{\boldsymbol{\pi}}$; the third inequality follows since $z(\tau, \tilde{\mathbf{u}}_0) \geq 0$ for all $\tau \geq 0$. Hence, the integral on the left hand side is uniformly bounded by a constant (independent of τ) for all $\tau \geq 0$. This implies that the integral must converge as $\tau \rightarrow \infty$. This completes the proof. \blacksquare

For the multi-class case, proving the uniqueness and global stability of the equilibrium point of the mean field is much more challenging since the monotonicity property, similar to the one established in Lemma 2, (for the single class case) does not hold in this case. We currently do not have a proof of global stability for the multi-class case. However, some case studies (as given below) strongly suggest that the uniqueness and global stability of the equilibrium point may be true in general.

1) *Case study 1:* Let us consider a multi-class system with two classes ($L = 2$), where the servers are identical with capacity $C = 3$ and the resource requirements for the two classes are $A_1 = 2$ and $A_2 = 3$, respectively. In this case, the state space of each server consists of three states given by $\mathcal{S} = \{(0, 0), (1, 0), (0, 1)\}$, where the first and the second component in each state respectively denote the numbers of jobs in the first and the second class that are in progress at the server. Let x_{00} , x_{10} , and x_{01} denote the fractions of servers in states $(0, 0)$, $(1, 0)$, and $(0, 1)$, respectively. We have from (12) that

$$\dot{x}_{00} = -(\lambda_1 + \lambda_2)(1 - (x_{10} + x_{01})^d) + x_{10} + x_{01}, \quad (29)$$

$$\dot{x}_{10} = \lambda_1(1 - (x_{10} + x_{01})^d) - x_{10}, \quad (30)$$

$$\dot{x}_{00} = \lambda_2(1 - (x_{10} + x_{01})^d) - x_{01}. \quad (31)$$

Now let $\pi = (\pi_{00}, \pi_{10}, \pi_{01})$ be an equilibrium point of the above system and define a *potential function* $V : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ as

$$V(t) = \frac{1}{\lambda_1}(x_{10} - \pi_{10})^2 + \frac{1}{\lambda_2}(x_{01} - \pi_{01})^2 \quad (32)$$

Taking the derivative with respect to time and using (29), (30), and (31) we obtain

$$\begin{aligned} \dot{V}(t) &= -\frac{2}{\lambda_1}(x_{(1,0),1}(t) - \pi_{(1,0),1})^2 - \frac{2}{\lambda_2}(x_{(0,1),1}(t) - \pi_{(0,1),1})^2 \\ &\quad - (x_{(1,0),1}(t) - \pi_{(1,0),1} + x_{(0,1),1}(t) - \pi_{(0,1),1})^2 \\ &\quad \times \left(\sum_{r=0}^{d-1} (\pi_{(1,0),1} + \pi_{(0,1),1})^r (x_{(1,0),1}(t) + x_{(0,1),1}(t))^{d-1-r} \right) \\ &\leq 0 \end{aligned} \quad (33)$$

This proves that global asymptotic stability (and hence uniqueness) of the equilibrium point of the mean field in the case under consideration.

2) *Case study 2:* Let us consider a multi-class system with two classes ($L = 2$), where the servers are identical with capacity $C = 4$ and the resource requirements for the two classes are $A_1 = 2$ and $A_2 = 3$, respectively. In this case, the state space of each server consists of four states given by $\mathcal{S} = \{(0, 0), (1, 0), (0, 1), (2, 0)\}$. Let x_{00} , x_{10} , x_{01} , and x_{20} denote the fractions

of servers in states $(0, 0)$, $(1, 0)$, $(0, 1)$, and $(2, 0)$ in the limiting system, respectively. Thus, from (12) we have for $d = 2$ that

$$\dot{x}_{00} = h_{00}(\mathbf{x}) = -(\lambda_1 + \lambda_2)(1 - (x_{10} + x_{01} + x_{20})^2) + x_{10} + x_{01} \quad (34)$$

$$\dot{x}_{10} = h_{10}(\mathbf{x}) = \lambda_1(1 - 2(x_{10} + x_{01} + x_{20})^2 + (x_{01} + x_{20})^2) - x_{10} + 2x_{20} \quad (35)$$

$$\dot{x}_{01} = h_{01}(\mathbf{x}) = \lambda_2(1 - (x_{10} + x_{01} + x_{20})^2) - x_{01} \quad (36)$$

$$\dot{x}_{20} = h_{20}(\mathbf{x}) = \lambda_1((x_{10} + x_{01} + x_{20})^2 - (x_{01} + x_{20})^2) - 2x_{20} \quad (37)$$

It is sufficient to consider only the last three components, i.e., $\mathbf{x} = (x_{10}, x_{01}, x_{20})$ to describe the system. It is clear that if the system is started at $\mathbf{x}(0) \in [0, 1]^3$, then it always remains in $[0, 1]^3$. We now calculate the Jacobian $J(\mathbf{x})$ of the $\mathbf{h}(\mathbf{x}) = (h_{10}(\mathbf{x}), h_{01}(\mathbf{x}), h_{20}(\mathbf{x}))$ at \mathbf{x} as follows

$$J(\mathbf{x}) = \begin{bmatrix} -4\lambda_1(x_{10} + x_{01} + x_{20}) - 1 & -2\lambda_1(2x_{10} + x_{01} + x_{20}) & -2\lambda_1(2x_{10} + x_{01} + x_{20}) + 2 \\ -2\lambda_2(x_{10} + x_{01} + x_{20}) & -2\lambda_2(x_{10} + x_{01} + x_{20}) - 1 & -2\lambda_2(x_{10} + x_{01} + x_{20}) \\ 2\lambda_1(x_{10} + x_{01} + x_{20}) & 2\lambda_1 x_{10} & 2\lambda_1 x_{10} - 2 \end{bmatrix}. \quad (38)$$

The characteristic polynomial of the above Jacobian matrix is given by $\Phi(z) = z^3 + a_2 z^2 + a_1 z + a_0$, where

$$a_2 = 4 + 2(x_{10} + 2(x_{01} + x_{20}))\lambda_1 + 2(x_{10} + x_{01} + x_{20})\lambda_2 \quad (39)$$

$$a_1 = 5 + 4(x_{10} + 2(x_{01} + x_{20}))\lambda_1 + 4(x_{01} + x_{20})(x_{10} + x_{01} + x_{20})\lambda_1^2 + 6(x_{10} + x_{01} + x_{20})\lambda_2 \\ + 4(x_{01} + x_{20})(x_{10} + x_{01} + x_{20})\lambda_1\lambda_2 \quad (40)$$

$$a_0 = 2 + 2(x_{10} + 2(x_{01} + x_{20}))\lambda_1 + 4(x_{01} + x_{20})(x_{10} + x_{01} + x_{20})\lambda_1^2 + 4(x_{10} + x_{01} + x_{20})\lambda_2 \quad (41)$$




Fig. 1. Squared Euclidean distance between $\mathbf{x}(t)$ and the equilibrium point $\boldsymbol{\pi}$ as a function of time for $\lambda_1 = 10$ and $\lambda_2 = 15$.

We observe that all the coefficients of this polynomial have the same sign. Furthermore, we have

$$\begin{aligned}
 a_2 a_1 - a_0 &= 18 + 4\lambda_1(6(x_{10} + 2(x_{01} + x_{20}))) + (2x_{10}^2 + 11x_{10}(x_{01} + x_{20})) \\
 &\quad + 11(x_{01} + x_{20})^2\lambda_1 + 2(x_{01} + x_{20})(x_{10} + x_{01} + x_{20})(x_{10} + 2(x_{01} + x_{20}))\lambda_1^2 \\
 &\quad + 2(x_{10} + x_{01} + x_{20})(15 + 2\lambda_1(2(x_{01} + x_{20}))(7 + 3(x_{01} + x_{20})\lambda_1) + x_{10}(5 + 4(x_{01} + x_{20})\lambda_1)))\lambda_2 + \\
 &\quad 2(x_{10} + x_{01} + x_{20})^2(3 + 4(x_{01} + x_{20})\lambda_1)\lambda_2^2 > 0, \quad (42)
 \end{aligned}$$

Hence, by the Routh array test we conclude that $J(\mathbf{x})$ is Hurwitz (hence invertible) for all $\mathbf{x} \in [0, 1]^3$. We also have that $\mathbf{h}(\mathbf{x}) = (h_{10}(\mathbf{x}), h_{01}(\mathbf{x}), h_{20}(\mathbf{x}))$ is a proper mapping (i.e., $\|\mathbf{h}(\mathbf{x})\| \rightarrow \infty$ as $\|\mathbf{x}\| \rightarrow \infty$) from \mathbb{R}^3 to \mathbb{R}^3 . The above two facts imply uniqueness of the equilibrium point of the system according to Hadamard's global inverse function theorem. Let $\boldsymbol{\pi}$ denote the equilibrium point.

We simulate the dynamics of the system with $\lambda_1 = 10$ and $\lambda_2 = 15$ and with random initial points in $[0, 1]^3$. In Figure 1, we plot the squared Euclidean distance between the $\mathbf{x}(t)$ and the equilibrium point $\boldsymbol{\pi}$ as a function of time. We observe that for all initial points, the trajectories converge asymptotically to the equilibrium point.

We note that for each N the process $\mathbf{x}^{(N)}(\cdot)$ is positive recurrent and hence has a unique stationary distribution π_N . We denote by $\mathbf{x}^{(N)}(\infty)$ the random variable distributed according to π_N . In the next theorem, we show that if the equilibrium point of the mean field is globally asymptotically stable then $\mathbf{x}^{(N)}(\infty)$ concentrates near the unique equilibrium point π of the mean field as $N \rightarrow \infty$.

Theorem 4: *Let π_N denote the stationary distribution of the process $\mathbf{x}^{(N)}(\cdot)$. If π is the globally asymptotically stable equilibrium point of the mean field, then the sequence $(\pi_N)_N$ converges weakly to δ_π as $N \rightarrow \infty$ or equivalently $\mathbf{x}^{(N)}(\infty) \Rightarrow \pi$*

Proof: We note that since the space \mathcal{U} is compact, the sequence of probability measures on $(\pi_N)_N$ is tight. Hence, Prohorov's theorem [22] implies that the sequence $(\pi_N)_N$ has limit points. We now show that all the limit points coincide with π . Theorem 1 implies that any limit point of the sequence $(\pi_N)_N$ is an invariant measure of the map $\mathbf{u}_0 \mapsto \mathbf{x}(t, \mathbf{u}_0)$. The global asymptotic stability of the equilibrium point implies that the unique invariant measure of the map $\mathbf{u}_0 \mapsto \mathbf{x}(t, \mathbf{u}_0)$ is δ_π . Therefore, any convergent subsequence of $(\pi_N)_N$ converges to δ_π . This completes the proof. ■

V. PROPAGATION OF CHAOS

In this section, we focus on the states of a given finite set of servers as $N \rightarrow \infty$. We show that as the system size grows the server occupancies become independent of each other which is formally known as the *propagation of chaos* or *asymptotic independence property*. We further show that if the equilibrium point of the mean field is globally stable, then the stationary distribution of server states in the limiting system is determined by the equilibrium point π of the mean field system (10)-(11). To formally state the results we introduce the following notations.

- The state of the k^{th} server of type j at a finite time $t \geq 0$ and at equilibrium are respectively denoted by the random variables $q_{k,j}^{(N)}(t)$ and $q_{k,j}^{(N)}(\infty)$, for $k \in \{1, 2, \dots, N\gamma_j\}$, $j \in \mathcal{J}$.
- For each $j \in \mathcal{J}$ and $t \geq 0$, we denote by $x_j(t, \mathbf{u}_0)$, the distribution on \mathcal{S}_j given by $x_j(t, \mathbf{u}_0) = (x_{\underline{n},j}(t, \mathbf{u}_0), n \in \mathcal{S}_j)$, where $\mathbf{x}(t, \mathbf{u}_0)$ is the mean field limit starting from \mathbf{u}_0 . Further, we define $x_j(\infty) = (\pi_{\underline{n},j}, n \in \mathcal{S}_j)$.

Further, we define the following notion of exchangeable random variables.

Definition *Let $\{q_{k,j}^{(N)}, 1 \leq k \leq N\gamma_j, 1 \leq j \leq M\}$ denote a collection of N random variables classified into M different types. The collection is called *intra-type exchangeable* if the joint law*

of the collection is invariant under permutation of indices, $1 \leq k \leq N\gamma_j$, of random variables belonging to type j for each $j \in \{1, 2, \dots, M\}$.

Theorem 5: For the model considered in this paper, if the equilibrium point π is globally asymptotically stable, $\{q_{k,j}^{(N)}(0), 1 \leq k \leq N\gamma_j, 1 \leq j \leq M\}$ is intra-type exchangeable, and $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \mathcal{U}$ as $N \rightarrow \infty$, then the following hold

- 1) For each fix k and $t \in [0, \infty]$, $\mathcal{L}(q_{k,j}^{(N)}(t)) \Rightarrow x_j(t, \mathbf{u}_0)$ as $N \rightarrow \infty$.
- 2) Fix positive integers r_1, r_2, \dots, r_M . For each $t \in [0, \infty]$,

$$\left\{ q_{k,j}^{(N)}(t), 1 \leq k \leq r_j, 1 \leq j \leq M \right\} \Rightarrow \{U_{k,j}(t), 1 \leq k \leq r_j, 1 \leq j \leq M\}, \quad (43)$$

as $N \rightarrow \infty$, where $U_{k,j}(t)$, $1 \leq k \leq r_j, 1 \leq j \leq M$, are independent random variables with $U_{k,j}(t)$ having distribution $x_j(t, \mathbf{u}_0)$ for all $1 \leq k \leq r_j$.

Proof: Note that the first part of Theorem 5 is a special case of the second part. Hence, it is sufficient to prove the second part. We will provide a proof for the $M = 2$ case. The proof readily extends to any $M \geq 2$.

Due to the dynamics of the system (power-of- d scheme) the joint law of the collection $\{q_{k,j}^{(N)}(t), 1 \leq k \leq N\gamma_j, 1 \leq j \leq 2\}$ depends only on the empirical distribution of states at time t given by $\mathbf{x}^{(N)}(t)$. Hence, permuting states among servers of the same type does not affect the joint law of the collection. Therefore, $\{q_{k,j}^{(N)}(t), 1 \leq k \leq N\gamma_j, 1 \leq j \leq 2\}$ is intra-type exchangeable for all $t \in [0, \infty]$. Now, given that $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \mathcal{U}$ as $N \rightarrow \infty$ we know from Theorem 1 and Theorem 4 that $\mathbf{x}^{(N)}(t) \Rightarrow \mathbf{x}(t, \mathbf{u}_0)$ as $N \rightarrow \infty$ for all $t \in [0, \infty]$. Henceforth, we will omit the variables t and \mathbf{u}_0 in our calculations since they hold for all $t \in [0, \infty]$ and all $\mathbf{u}_0 \in \mathcal{U}$. To prove the independence, it is sufficient to show that the following holds:

$$\mathbb{E} \left[\prod_{k=1}^{r_1} \phi_k \left(q_{k,1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{k,2}^{(N)} \right) \right] \rightarrow \prod_{k=1}^{r_1} \langle \phi_k, x_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2 \rangle \text{ as } N \rightarrow \infty \quad (44)$$

for all bounded mappings $\phi_k : \mathcal{S}_1 \rightarrow \mathbb{R}_+$ and $\psi_k : \mathcal{S}_2 \rightarrow \mathbb{R}_+$. We have

$$\begin{aligned}
& \left| \mathbb{E} \left[\prod_{k=1}^{r_1} \phi_k \left(q_{k,1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{k,2}^{(N)} \right) \right] - \prod_{k=1}^{r_1} \langle \phi_k, x_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2 \rangle \right| \\
& \leq \left| \mathbb{E} \left[\prod_{k=1}^{r_1} \phi_k \left(q_{k,1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{k,2}^{(N)} \right) \right] - \mathbb{E} \left[\prod_{k=1}^{r_1} \langle \phi_k, x_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2^{(N)} \rangle \right] \right| \\
& \quad + \left| \mathbb{E} \left[\prod_{k=1}^{r_1} \langle \phi_k, x_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2^{(N)} \rangle \right] - \prod_{k=1}^{r_1} \langle \phi_k, x_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2 \rangle \right|, \quad (45)
\end{aligned}$$

where $x_j^{(N)} = (x_{\underline{n},j}, \underline{n} \in \mathcal{S}_j)$ is the random probability measure on \mathcal{S}_j induced by the process $\mathbf{x}^{(N)}$. We note that the second term on the right hand side of the above inequality vanishes as $N \rightarrow \infty$ because of the following facts: $x_j^{(N)} \Rightarrow x_j$ as $N \rightarrow \infty$ for $j = 1, 2$; x_1 and x_2 are deterministic; $x_j^{(N)}$ is a bounded random vector for $j = 1, 2$. Now, due to intra-type exchangeability the permutation of states between servers belonging to the same class does not affect the joint distribution. Hence, we have

$$\begin{aligned}
\mathbb{E} \left[\prod_{k=1}^{r_1} \phi_k \left(q_{k,1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{k,2}^{(N)} \right) \right] &= \frac{1}{(N\gamma_1)^{r_1} (N\gamma_2)^{r_2}} \times \\
& \mathbb{E} \left[\sum_{\substack{\sigma \in P(r_1, N\gamma_1) \\ \sigma' \in P(r_2, N\gamma_2)}} \prod_{k=1}^{r_1} \phi_k \left(q_{\sigma(k),1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{\sigma'(k),2}^{(N)} \right) \right] \quad (46)
\end{aligned}$$

where $(N)_k = N(N-1)\dots(N-k+1)$, and $P(r, n)$ denotes the set of all permutations of the numbers $\{1, 2, \dots, n\}$ taken r at a time. Also, by definition of $x_j^{(N)}$ we have

$$\begin{aligned}
& \mathbb{E} \left[\prod_{k=1}^{r_1} \langle \phi_k, x_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2^{(N)} \rangle \right] \\
& = \mathbb{E} \left[\left(\prod_{k=1}^{r_1} \frac{1}{N\gamma_1} \sum_{l=1}^{N\gamma_1} \phi_k \left(q_{l,1}^{(N)} \right) \right) \left(\prod_{k=1}^{r_2} \frac{1}{N\gamma_2} \sum_{l=1}^{N\gamma_2} \psi_k \left(q_{l,2}^{(N)} \right) \right) \right] \quad (47)
\end{aligned}$$

Hence, the first term on the right hand side of (45) can be bounded as follows

$$\begin{aligned}
& \left| \mathbb{E} \left[\prod_{k=1}^{r_1} \phi_k \left(q_{k,1}^{(N)} \right) \prod_{k=1}^{r_2} \psi_k \left(q_{2,k}^{(N)} \right) \right] - \mathbb{E} \left[\prod_{k=1}^{r_1} \langle \phi_k, x_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, x_2^{(N)} \rangle \right] \right| \\
& \leq (N\gamma_1)_{r_1} (N\gamma_2)_{r_2} \left(\frac{1}{(N\gamma_1)_{r_1} (N\gamma_2)_{r_2}} - \frac{1}{(N\gamma_1)^{r_1} (N\gamma_2)^{r_2}} \right) B^{r_1+r_2} \\
& \quad + ((N\gamma_1)^{r_1} (N\gamma_2)^{r_2} - (N\gamma_1)_{r_1} (N\gamma_2)_{r_2}) \frac{B^{r_1+r_2}}{(N\gamma_1)^{r_1} (N\gamma_2)^{r_2}} \\
& \leq 2B^{r_1+r_2} \left(1 - \frac{(N\gamma_1)_{r_1} (N\gamma_2)_{r_2}}{(N\gamma_1)^{r_1} (N\gamma_2)^{r_2}} \right) \rightarrow 0 \text{ as } N \rightarrow \infty
\end{aligned}$$

where B is a constant such that $\|\phi_k\|_\infty < B$ for $k = 1, 2, \dots, r_1$ and $\|\psi_k\|_\infty < B$ for $k = 1, 2, \dots, r_2$. This completes the proof. \blacksquare

Remark 4: We note that any initial condition that specifies $\mathbf{x}_{\underline{n},j}^{(N)}(0) = g_{\underline{n},j}$ for all N , $\underline{n} \in \mathcal{S}_j$, and $j \in \mathcal{J}$, satisfies the conditions in Theorem 5. In particular, the conditions in Theorem 5 are satisfied when all the servers are empty at $t = 0$.

Thus, the above theorem shows that in the limiting system any finite set of servers become independent of each other and the stationary distribution of states of any server of type- j is given by $\pi_j = \{\pi_{\underline{n},j}, \underline{n} \in \mathcal{S}_j\}$. The following proposition shows (using the independence of the servers in the limiting system) that in equilibrium the arrivals of class $l \in \mathcal{L}$ at any given server of type $j \in \mathcal{J}$ in the limiting system form a state dependent Poisson process whose rates are given by $\lambda_{\underline{n},j}(\boldsymbol{\pi})$, $\underline{n} \in \mathcal{S}_j$, where $\lambda_{\underline{n},j}(\mathbf{x})$ for $\mathbf{x} \in \mathcal{U}$ are as defined in (14).

Proposition 2: In equilibrium, the arrival process of jobs at any given server in the limiting system is a state dependent Poisson process. Furthermore, the equilibrium arrival rate of class- l jobs at a server of type $j \in \mathcal{J}$, when it is in state $\underline{n} \in \mathcal{S}_j$, is given by

$$\lambda_{\underline{n},j}^{(l)}(\boldsymbol{\pi}) = \lambda_l \frac{F(\underline{n}, j, \boldsymbol{\pi})}{E(\underline{n}, j, j, \boldsymbol{\pi})}, \quad (48)$$

where $\underline{n} \in \mathcal{S}_j$ is such that $\underline{n} + \underline{e}_l \in \mathcal{S}_j$ and $F(\underline{n}, j, \boldsymbol{\pi})$, $E(\underline{n}, j, j, \boldsymbol{\pi})$ are as defined in Lemma 1.

Proof: The proof is given in the Appendix D. \blacksquare

From Remark 3, we already know that the equilibrium point $\boldsymbol{\pi}$ of the mean field satisfies

$$\lambda_{\underline{n}-\underline{e}_l,j}^{(l)}(\boldsymbol{\pi}) \pi_{\underline{n}-\underline{e}_l,j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = n_l \pi_{\underline{n},j} \text{ for } \underline{n} \in \mathcal{S}_j \text{ and } l \in \mathcal{L}. \quad (49)$$

Now, since (by Proposition 2) $\lambda_{\underline{n}-\underline{e}_l,j}^{(l)}(\boldsymbol{\pi})$ is the equilibrium arrival rate of class- l jobs at a server of type j in state $\underline{n} - \underline{e}_l \in \mathcal{S}_j$, the above equations can be interpreted as the detailed balance

equations that equate the transition rates between the states $\underline{n} - \underline{e}_l$ and \underline{n} for each $\underline{n}, \underline{n} - \underline{e}_l \in \mathcal{S}_j$, $j \in \mathcal{J}$, $l \in \mathcal{L}$. Using the detailed balance equations, we now find a recursive relationship among the stationary tail probabilities of the number of occupied resource units as in [23], [24] at each server in the limiting system. This allows efficient computation of the blocking probabilities for each class of jobs.

Proposition 3: Let $P_{k,j}$, for $1 \leq k \leq C_j$ and $j \in \mathcal{J}$, denote the stationary probability that a server in the limiting system has at least k units of occupied resources, i.e., $P_{k,j} = \sum_{\substack{\underline{n} \in \mathcal{S}_j: \\ \underline{n} \cdot \underline{A} \geq k}} \pi_{\underline{n},j}$. Then $P_{k,j}$ satisfies (1) for $0 \leq k \leq C_j - 1$, where $P_{k,j} = 1$ for $k \leq 0$, and $P_{C_j+1,j} = 0$ for all $j \in \mathcal{J}$.

Proof: For $j \in \mathcal{J}$, $0 \leq k \leq C_j$, we define the set $\mathcal{D}_{k,j}$ as $\mathcal{D}_{k,j} = \{\underline{n} \in \mathcal{S}_j : \underline{n} \cdot \underline{A} = k\}$. Thus, $\mathcal{D}_{k,j}$ denotes the set of states in \mathcal{S}_j for which the total number occupied VM's at a type j server is exactly k . We note that for all $\underline{n} \in \mathcal{D}_{k,j}$ such that $\underline{n} - \underline{e}_l \in \mathcal{S}_j$, we have $G(\underline{n} - \underline{e}_l, i, j, \boldsymbol{\pi}) = \gamma_i P_{k-A_l+C_i-C_j+1,i}$ and $E(\underline{n} - \underline{e}_l, i, j, \boldsymbol{\pi}) = \gamma_i (P_{k-A_l+C_i-C_j,i} - P_{k-A_l+C_i-C_j+1,i})$. Thus, for all $\underline{n} \in \mathcal{D}_{k,j}$ such that $\underline{n} - \underline{e}_l \in \mathcal{S}_j$ we have

$$\lambda_{\underline{n} - \underline{e}_l, j}^{(l)} = \frac{\lambda_l}{\gamma_j (P_{k-A_l, j} - P_{k-A_l, j})} \left[\left(\sum_{i=1}^j \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j+1}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d \right]. \quad (50)$$

Now from (49) we have

$$\sum_{l \in \mathcal{L}} \sum_{\underline{n} \in \mathcal{D}_{k,j}} A_l \lambda_{\underline{n} - \underline{e}_l, j}^{(l)} \pi_{\underline{n} - \underline{e}_l, j} I_{\underline{n} - \underline{e}_l \in \mathcal{S}_j} = \sum_{l \in \mathcal{L}} \sum_{\underline{n} \in \mathcal{D}_{k,j}} n_l A_l \pi_{\underline{n}, j} \text{ for } \underline{n} \in \mathcal{S}^j \text{ and } l \in \mathcal{L} \quad (51)$$

Now, the LHS of the above equation can be simplified as follows:

$$\begin{aligned}
\sum_{l \in \mathcal{L}} \sum_{\underline{n} \in \mathcal{D}_{k,j}} A_l \lambda_{\underline{n}-\underline{e}_l, j}^{(l)} \pi_{\underline{n}-\underline{e}_l, j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} &= \sum_{l \in \mathcal{L}} A_l \lambda_{\underline{n}-\underline{e}_l, j}^{(l)} \sum_{\underline{n} \in \mathcal{D}_{k,j}} \pi_{\underline{n}-\underline{e}_l, j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} \\
&= \sum_{l \in \mathcal{L}} A_l \frac{\lambda_l}{\gamma_j (P_{k-A_l, j} - P_{k-A_l+1, j})} \left[\left(\sum_{i=1}^j \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j+1}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d \right. \\
&\quad \left. - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d \right] (P_{k-A_l, j} - P_{k-A_l+1, j}) \\
&= \sum_{l \in \mathcal{L}} A_l \frac{\lambda_l}{\gamma_j} \left[\left(\sum_{i=1}^j \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j+1}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d \right. \\
&\quad \left. - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-A_l+C_i-C_j, i} + \sum_{i=j}^M \gamma_i P_{k-A_l+C_i-C_j+1, i} \right)^d \right]
\end{aligned}$$

The second equality follows since $\sum_{\underline{n} \in \mathcal{D}_{k,j}} \pi_{\underline{n}-\underline{e}_l, j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = (P_{k-A_l, j} - P_{k-A_l+1, j})$. Similarly, the RHS can be simplified as

$$\sum_{l \in \mathcal{L}} \sum_{\underline{n} \in \mathcal{D}_{k,j}} n_l A_l \pi_{\underline{n}, j} = \sum_{\underline{n} \in \mathcal{D}_{k,j}} \pi_{\underline{n}, j} \sum_{l \in \mathcal{L}} n_l A_l = \sum_{\underline{n} \in \mathcal{D}_{k,j}} \pi_{\underline{n}, j} k = k (P_{k, j} - P_{k+1, j}).$$

This completes the proof. ■

Remark 5 (Insensitivity): All the results, discussed so far in this section, have been obtained assuming that the service time distribution of the incoming jobs is exponential. The same results can be shown to hold for any service time distribution if asymptotic independence of the servers is assumed to hold for general service time distributions. The asymptotic independence property was conjectured to hold for homogeneous systems with local service disciplines and general service time distributions in [25]. The proof of this remains as an open problem.

Under the assumption of asymptotic independence of servers for general service time distributions, the statement of Proposition 2 continues to hold, i.e., the equilibrium arrival process at each server in the limiting system is a state dependent Poisson process whose rates are given by (48). This implies that the detailed balance equations given by (49) also hold for general service time distributions. Since the servers in the system are loss servers, the detailed balance condition implies that the stationary distribution of each server in the limiting system is *insensitive* to service time distributions (see Theorem 1 of [26]). We refer to this property as the




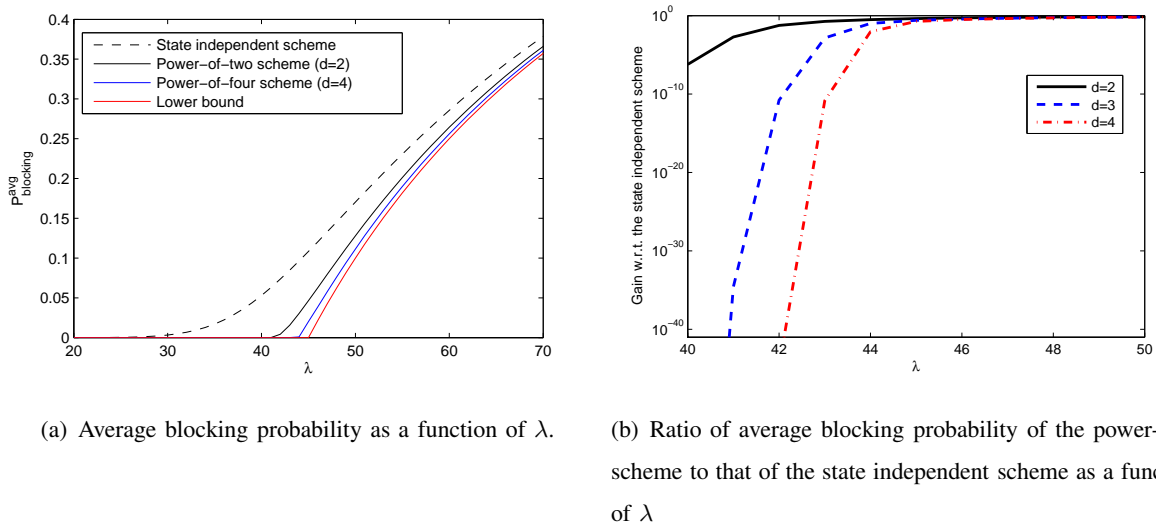
Fig. 2. Accuracy of mean field analysis of power-of- d scheme: Average blocking probability as a function of λ for different values of N .

asymptotic insensitivity of the system. Thus, asymptotic insensitivity of the system holds under the hypothesis of asymptotic independence of the servers, the proof of which remains as an open problem. In the next section, we provide numerical evidences to support insensitivity.

VI. NUMERICAL RESULTS

We first investigate the accuracy of the asymptotic analysis presented in the paper in predicting the system performance for finite system size N . We set the following parameter values: $L = 1$, $A_1 = A = 1$, $M = 2$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 20$, $C_2 = 25$, and $d = 2$. All simulation results presented in this section are the average of 10,000 independent runs. We use λ to denote the arrival rate of jobs. In Figure 2, we plot the average blocking probability of requests under the power-of-two scheme as a function of λ for $N = 10, 50$. We have also plotted the blocking probability obtained by solving (3). We observe that results obtained from the simulations match almost exactly with those obtained from the analysis. This leads us to believe that the mean-field results derived in this paper very accurately predict the behavior of the power-of- d scheme even for moderate system sizes.

In Figure 3 we compare for $L = 1$ the average blocking probability of the power-of- d scheme with that of the state independent routing scheme, in which an incoming job is routed to a server of type j with probability p_j independent of the states of the servers. We set $p_j =$

Fig. 3. Efficacy of the power-of- d scheme

$\gamma_j C_j / \sum_{i=1}^M \gamma_i C_i$ since in this case p_j is proportional to both γ_j and C_j .¹ The parameters were chosen as $M = 2$, $N = 100$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 30$, and $C_2 = 60$. For this parameter setting, the critical load is given by $\lambda_{\text{crit}} = \gamma_1 C_1 + \gamma_2 C_2 = 45$. In Figure 3(a), we plot the average blocking probability of the two schemes as a function of the arrival rate λ . We have also plotted the lower bound obtained from (2). In Figure 3(b), we plot the ratio of average blocking probability of the power-of- d scheme to that of the state independent scheme. Note that the y-axis is in the log scale. From Figure 3(a) we observe that the average blocking probability obtained for $d = 4$ is almost equal to that of the lower bound. We also observe from Figure 3(b), that the average blocking probability under the power-of- d scheme is orders of magnitude lower than that under the state independent routing scheme around $\lambda = \lambda_{\text{cap}}$. This shows the efficacy of such randomized strategies in reducing blocking for realistic systems which are typically operated near the critical load.

We now numerically confirm the insensitivity of the power-of-two scheme under different service time distributions. We set the following parameter values: $M = 2$, $d = 2$, $N = 100$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 20$, and $C_2 = 25$. In Table I, average blocking probability is shown as a function of λ , for the following distributions: 1) *Constant*: We consider job length distribution

¹The probabilities p_j , $j \in \mathcal{J}$, can be optimally chosen to minimize the average blocking probability. However, such optimal choice requires the knowledge of the arrival rate λ , which is difficult to estimate.

having the cumulative distribution given by $F(x) = 0$ for $0 \leq x < 1$, and $F(x) = 1$, otherwise and 2) *Power law*: We consider job length distribution having cumulative distribution function given by $F(x) = 1 - 1/4x^2$ for $x \geq \frac{1}{2}$ and $F(x) = 0$, otherwise. Note that for each of the above distributions the average service time is 1. We see from Table I that the change in blocking probability is insignificant when the service time distribution is changed keeping the same mean. This supports the fact the under the power-of- d scheme the system is insensitive to the service time distribution in the limit as $N \rightarrow \infty$ (asymptotic insensitivity).

TABLE I
ASYMPTOTIC INSENSITIVITY OF THE POWER-OF- d SCHEME

λ	Constant (Simulation)	Power Law (Simulation)
20	0.0087	0.0086
25	0.1467	0.1470
30	0.2758	0.2747
35	0.3733	0.3737
40	0.4490	0.4485
45	0.5085	0.5085

VII. CONCLUDING REMARKS

In this paper, we analyzed the power-of- d scheme for multi-class heterogeneous Erlang loss systems with a large number of servers. We showed that in the large system limit the evolution of the empirical occupancy distribution can be characterized through its mean field limit. Furthermore, we showed that propagation of chaos holds for heterogeneous case through the requirement of intra-type exchangeability.

VIII. ACKNOWLEDGEMENTS

We thank Prof. R. Srikant for pointing out a gap in the original version of the paper pertaining to the proof of uniqueness and asymptotic convergence to the fixed point. This pertains to Theorems 3 and 4 in this version. We also wish to thank Prof. C. Nielsen for pointing us in the direction of the Krasovskii-Lyapunov theorem.

REFERENCES

- [1] “Amazon EC2.” <http://aws.amazon.com/ec2/>.
- [2] “Microsoft Azure.” <http://www.microsoft.com/windowsazure/>.
- [3] V. Gupta and A. Radovanovic, “Online stochastic bin packing.” *Corr: abs/1211.2687*, 2012.
- [4] Q. Xie, X. Dong, Y. Lu, and R. Srikant, “Power of d choices for large-scale bin packing: A loss model,” in *Proceedings of ACM SIGMETRICS 2015*.
- [5] S. T. Maguluri, R. Srikant, and L. Ying, “Stochastic models of load balancing and scheduling in cloud computing clusters,” in *Proceedings of IEEE INFOCOM*, 2012.
- [6] S. T. Maguluri, R. Srikant, and L. Ying, “Heavy traffic optimal resource allocation algorithms for cloud computing clusters,” in *Proceedings of 24th International Teletraffic Congress (ITC 24)*, 2012.
- [7] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich, “Queueing system with selection of the shortest of two queues: an asymptotic approach,” *Problems of Information Transmission*, vol. 32, no. 1, pp. 20–34, 1996.
- [8] M. Mitzenmacher, “The power of two choices in randomized load balancing,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 12, no. 10, pp. 1094–1104, 2001.
- [9] A. Ganesh, S. Lilienthal, D. Manjunath, A. Proutiere, and F. Simatos, “Load balancing via random local search in closed and open systems,” *Queueing Systems*, 2012.
- [10] 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*, 2015, to appear.
- [11] S. R. E. Turner, “Resource pooling in stochastic networks,” *Ph.D. dissertation, University of Cambridge*, 1996.
- [12] S. R. E. Turner, “The effect of increasing routing choice on resource pooling,” *Probability in the Engineering and Informational Sciences*, vol. 12, pp. 109–124, 1998.
- [13] C. Graham and S. Méléard, “Propagation of chaos for a fully connected loss network with alternate routing,” *Stochastic Processes and their Applications*, vol. 44, no. 1, pp. 159–180, 1993.
- [14] C. Graham and S. Méléard, “Stochastic particle approximations for generalized Boltzmann models and convergence estimates,” *The Annals of Probability*, vol. 28, no. 1, pp. 115–132, 1997.
- [15] M. Mitzenmacher, “The power of two choices in randomized load balancing,” *PhD Thesis, Berkeley*, 1996.
- [16] J. B. Martin and Y. M. Suhov, “Fast Jackson networks,” *Annals of Applied Probability*, vol. 9, no. 3, pp. 854–870, 1999.
- [17] K. Deimling, “Ordinary differential equations in Banach spaces,” vol. 596 of *Lecture Notes in Mathematics*, Springer Berlin, 1977.
- [18] H. K. Khalil, *Nonlinear systems*. Macmillan Publishing Company, New York, 1992.
- [19] M. S. Berger, *Nonlinearity and Functional Analysis: Lectures on Nonlinear Problems in Mathematical Analysis*. Academic Press, 1977.
- [20] R. Hettich and K. O. Kortanek, “Semi-infinite programming: theory, methods, and applications,” *SIAM Rev.*, vol. 35, pp. 380–429, 1993.
- [21] B. Kuo, *Automatic Control Systems*. Prentice-Hall, 1982.
- [22] P. Billingsley, *Convergence of Probability Measures*. Wiley, 1968.
- [23] J. Kaufman, “Blocking in a shared resource environment,” *IEEE Transactions on Communications*, vol. 29, no. 10, pp. 1474–1481, 1981.
- [24] J. W. Roberts, “A service system with heterogeneous user requirement,” in *Performance of Data Communications Systems and Their Applications*, pp. 423–431, 1981.

- [25] 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.
- [26] S. Zachary, "A note on insensitivity in stochastic networks," *Journal of Applied Probability*, vol. 44, no. 1, pp. 238–248, 2007.
- [27] S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence*. John Wiley and Sons Ltd, 1985.

APPENDIX

A. Proof of Proposition 1

Using (3) it can be shown that the following recursive relation holds for $0 \leq k \leq C_M - 1$

$$\begin{aligned} \sum_{j \in \mathcal{J}} (k+1+C_j-C_M)_+ \gamma_j (P_{k+1+C_j-C_M,j} - P_{k+2+C_j-C_M,j}) \\ = \lambda \left[\left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right)^d - \left(\sum_{j=1}^M \gamma_j P_{k+1+C_j-C_M,j} \right)^d \right], \end{aligned} \quad (52)$$

where $(y)_+ = \max(0, y)$. From (52) the following can be shown to hold for $0 \leq k \leq C_M - 1$ using backward induction starting at $k = C_M - 1$.

$$\sum_{j=1}^M (k+1+C_j-C_M)_+ \gamma_j P_{k+1+C_j-C_M,j} \leq \lambda \left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right)^d. \quad (53)$$

From (53) it is clear that

$$\left(\sum_{j=1}^M \gamma_j P_{k+1+C_j-C_M,j} \right) \leq \frac{\lambda}{k+(C_1-C_M)+1} \left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right)^d, \quad (54)$$

for $C_M - C_1 \leq k \leq C_M - 1$. Now, for $0 \leq k \leq k_0$, we have $\bar{P}_k = 1 \geq \left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right)$. Assume that $\left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right) \leq \bar{P}_k$ holds for some $k \geq k_0$. Using induction, we will now show that the inequality must hold for $k+1$. We have

$$\begin{aligned} \left(\sum_{j=1}^M \gamma_j P_{k+1+C_j-C_M,j} \right) &\leq \frac{\lambda}{k+(C_1-C_M)+1} \left(\sum_{j=1}^M \gamma_j P_{k+C_j-C_M,j} \right)^d \\ &\leq \frac{\lambda}{k+(C_1-C_M)+1} \bar{P}_k^d = \bar{P}_{k+1}. \end{aligned}$$

This completes the proof. ■

B. Proof of Lemma 1

We first consider the transition of the system from the state $\mathbf{u} \in \mathcal{U}^{(N)}$ at t^- to the state $\mathbf{u} - \frac{\mathbf{e}(\underline{n}-\underline{e}_j)}{N\gamma_j} + \frac{\mathbf{e}(\underline{n},j)}{N\gamma_j}$ at t , where $\underline{n} \in \mathcal{S}_j$. This transition occurs when an arrival of class $l \in \mathcal{L}$

at time t joins a type- j server which was in state $\underline{n} - \underline{e}_l$ at time t^- (just before the arrival). Let k of the d sampled servers be of type j with state \underline{n}' satisfying $\underline{n}' \cdot \underline{A} = (\underline{n} - \underline{e}_l) \cdot \underline{A}$. For the transition to occur, we must have $k \geq 1$ and one among the k servers must be in state $\underline{n} - \underline{e}_l$. Since there are $N\gamma_j u_{\underline{n}-\underline{e}_l, j}$ and $N \times E(\underline{n} - \underline{e}_l, j, j, \mathbf{u})$ servers of type j in states $\underline{n} - \underline{e}_l$ and \underline{n}' satisfying $\underline{n}' \cdot \underline{A} = (\underline{n} - \underline{e}_l) \cdot \underline{A}$, respectively, the probability of sampling k such servers is $\binom{k}{1} \gamma_j u_{\underline{n}-\underline{e}_l, j} E^{k-1}(\underline{n} - \underline{e}_l, j, j, \mathbf{u})$. In this case, since there are k servers with equal vacancy, the arrival joins a server with state $\underline{n} - \underline{e}_l$ with probability $1/k$. The other $d - k$ sampled servers must satisfy either of the following two conditions:

- If the sampled server is of type $i < j$, then its state \underline{n}' must satisfy $C_i - \underline{n}' \cdot \underline{A} \leq C_j - (\underline{n} - \underline{e}_l) \cdot \underline{A}$, or, $\underline{n}' \cdot \underline{A} \geq (\underline{n} - \underline{e}_l) \cdot \underline{A} + C_i - C_j$. The number of type i servers in a state satisfying the above relation is $N \times GE(\underline{n} - \underline{e}_l, i, j, \mathbf{u})$. Since servers are sampled uniformly at random, the probability with which one of these servers is sampled is $GE(\underline{n} - \underline{e}_l, i, j, \mathbf{u})$.
- If the sampled server is of type $i \geq j$, then its state \underline{n}' must satisfy $C_i - \underline{n}' \cdot \underline{A} < C_j - (\underline{n} - \underline{e}_l) \cdot \underline{A}$, or, $\underline{n}' \cdot \underline{A} > (\underline{n} - \underline{e}_l) \cdot \underline{A} + C_i - C_j$. Using the similar argument as before, the probability with which such a server is sampled is $G(\underline{n} - \underline{e}_l, i, j, \mathbf{u})$.

Thus the total probability with which the incoming arrival joins a server of type j in state $\underline{n} - \underline{e}_l$ is $\sum_{k=1}^d \binom{k}{1} \frac{1}{k} \gamma_j u_{\underline{n}-\underline{e}_l, j} E^{k-1}(\underline{n} - \underline{e}_l, j, j, \mathbf{u}) \left(\sum_{i=1}^{j-1} GE(\underline{n}, i, j, \mathbf{u}) + \sum_{i=j}^M G(\underline{n}, i, j, \mathbf{u}) \right)^{d-k}$ which simplifies to $\frac{F(\underline{n}-\underline{e}_l, j, \mathbf{u})}{E(\underline{n}-\underline{e}_l, j, j, \mathbf{u})} \gamma_j u_{\underline{n}-\underline{e}_l, j}$. Since the arrival rate of class- l jobs is $N\lambda_l$, the rate of transition from the state \mathbf{u} to the state $\mathbf{u} - \frac{\mathbf{e}(\underline{n}-\underline{e}_l, j)}{N\gamma_j} + \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}$ is given by

$$r \left(\mathbf{u} \rightarrow \mathbf{u} - \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} + \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j} \right) = N\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \mathbf{u})}{E(\underline{n} - \underline{e}_l, j, j, \mathbf{u})} \gamma_j u_{\underline{n}-\underline{e}_l, j}. \quad (55)$$

Next, we consider the transition from the state $\mathbf{u} \in \mathcal{U}^{(N)}$ to the state $\mathbf{u} + \frac{\mathbf{e}(\underline{n}-\underline{e}_l, j)}{N\gamma_j} - \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}$, where $\underline{n} \in \mathcal{S}_j$. This transition occurs when a job of class $l \in \mathcal{L}$ leaves a type $j \in \mathcal{J}$ server in state \underline{n} . The number of type- j servers in state \underline{n} when the system is in state \mathbf{u} is $N\gamma_j u_{\underline{n}, j}$. From each of these servers, the rate at which class- l jobs depart is n_l . Hence, the rate of transition from the state \mathbf{u} to the state $\mathbf{u} + \frac{\mathbf{e}(\underline{n}-\underline{e}_l, j)}{N\gamma_j} - \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}$ is given by

$$r \left(\mathbf{u} \rightarrow \mathbf{u} + \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} - \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j} \right) = N\gamma_j u_{\underline{n}, j} n_l \quad (56)$$

The expression (5) now follows directly from the definition of $\mathbf{A}^{(N)}$.

C. Proof of Theorem 1

The proof consists of three main steps. The first step is to show that the sequence of Markov processes $\{\mathbf{x}^{(N)}(\cdot)\}_N$ is relatively compact. The second step is to show that there exists a unique process $\mathbf{x}(\cdot)$ satisfying (10)-(11). The third step is to show that the operator semigroup $(\mathbf{T}^{(N)}(t), t \geq 0)$ generated by $\mathbf{A}^{(N)}$ corresponding to the Markov process $\mathbf{x}^{(N)}(\cdot)$ converges to the operator semigroup of the process $\mathbf{x}(\cdot)$, i.e.,

$$\lim_{N \rightarrow \infty} \sup_{\mathbf{u} \in \mathcal{U}^{(N)}} |\mathbf{T}^{(N)}(t)f(\mathbf{u}) - f(\mathbf{x}(t, \mathbf{u}))| = 0, \quad (57)$$

for all continuous functions $f : \mathcal{U} \rightarrow \mathbb{R}$, where the convergence is uniform in t within any bounded interval. Combining these three steps, the statement of Theorem 1 follows from Corollary 8.7 of Chapter 4 of [27].

The proof of the first part is essentially the same as the proof of Theorem 6.1 of [11] but we give the details here for the completeness of the paper. We first recall that the metric ρ defined on the space \mathcal{U} is given by:

$$\rho(\mathbf{u}, \mathbf{w}) = \sup_{j \in \mathcal{J}} \sup_{\underline{n} \in \mathcal{S}_j} \frac{|u_{\underline{n}, j} - w_{\underline{n}, j}|}{(\underline{n} \cdot \underline{e}) + 1}. \quad (58)$$

Now, to prove relative compactness of the sequence of processes $\{\mathbf{x}^{(N)}(\cdot)\}$ we need to satisfy the following three conditions of Theorem 8.6 of Chapter 3 of [27]:

- For every $\eta > 0$ and rational $t \geq 0$, there exists a compact set $\Gamma_{\eta, t}$ such that

$$\liminf_{N \rightarrow \infty} \mathbb{P} [\inf \{\rho(\mathbf{x}^{(N)}(t), y) : y \in \Gamma_{\eta, t}\} < \eta] \geq 1 - \eta. \quad (59)$$

This is condition (7.7) of Chapter 3 of [27].

- For all $T > 0$, there exists $\beta > 0$, $C > 0$, and $\theta > 1$, such that for all N and all $0 \leq h \leq t \leq T + 1$,

$$\mathbb{E} [\rho^{\beta/2}(\mathbf{x}^{(N)}(t+h), \mathbf{x}^{(N)}(t)) \rho^{\beta/2}(\mathbf{x}^{(N)}(t), \mathbf{x}^{(N)}(t-h))] \leq Ch^\theta. \quad (60)$$

This is condition (8.37) of Chapter 3 of [27] which by Theorem 8.8 of Chapter 3 of [27] implies condition (8.28) of Theorem 8.6 of Chapter 3 of [27].

- For that β ,

$$\lim_{\delta \rightarrow 0} \limsup_{N \rightarrow \infty} \mathbb{E} [\rho^\beta(\mathbf{x}^{(N)}(\delta), \mathbf{x}^{(N)}(0))] = 0. \quad (61)$$

This is condition (8.30) of Chapter 3 of [27].

Now, (59) follows immediately by choosing $\Gamma_{\eta,t} = \mathcal{U}$ for all $\eta > 0$ and $t \geq 0$ since \mathcal{U} is compact and the process $\mathbf{x}^{(N)}(t)$ lies in the space $\mathcal{U}^{(N)} \subseteq \mathcal{U}$. Clearly, the change under the metric ρ is bounded above by that under the sup metric. Due to arrival or departure one of $x_{\underline{n},j}^{(N)}$ is increased by $1/N\gamma_j$ and one is decreased by $1/N\gamma_j$. This changes \mathbf{x}^N by at most $1/(N(\min_{j \in \mathcal{J}} \gamma_j))$ in the sup metric. We next need to bound the number of arrivals and departures that can occur in an interval of length h . The arrivals form a Poisson process whose rate is bounded above by $N\lambda$, where $\lambda = \sum_{l \in \mathcal{L}} \lambda_l$ and the number of departures is bounded above by the number of events in a Poisson process of rate NB where $B = \max_{j \in \mathcal{J}} \max_{\underline{n} \in \mathcal{S}_j} (\underline{n} \cdot \underline{e})$. Therefore, the total number of events in an interval of length h is bounded above by a Poisson random variable with mean $(\lambda + B)Nh$. Thus, $\rho(\mathbf{x}^{(N)}(t+h), \mathbf{x}^{(N)}(t))$ is bounded above by $1/(N(\min_{j \in \mathcal{J}} \gamma_j))$ times a Poisson random variable with mean $(\lambda + B)Nh$. Now we set $\beta = 2$. Then using the Markov property of $\mathbf{x}^{(N)}$ we have

$$\begin{aligned} & \mathbb{E} [\rho^{\beta/2}(\mathbf{x}^{(N)}(t+h), \mathbf{x}^{(N)}(t)) \rho^{\beta/2}(\mathbf{x}^{(N)}(t), \mathbf{x}^{(N)}(t-h))] \\ &= \mathbb{E} [\rho^{\beta/2}(\mathbf{x}^{(N)}(t+h), \mathbf{x}^{(N)}(t))] \mathbb{E} [\rho^{\beta/2}(\mathbf{x}^{(N)}(t), \mathbf{x}^{(N)}(t-h))] \leq \frac{(\lambda + B)^2}{(\min_{j \in \mathcal{J}} \gamma_j)^2} h^2. \end{aligned} \quad (62)$$

Hence, (60) holds for all necessary T, t , and h with $C = \frac{(\lambda+B)^2}{(\min_{j \in \mathcal{J}} \gamma_j)^2}$ and $\theta = 2$. It remains to prove (61) for $\beta = 2$. This is easy to see since $\mathbb{E} [\rho^\beta(\mathbf{x}^{(N)}(\delta), \mathbf{x}^{(N)}(0))] \leq \frac{(\lambda+B)N\delta + (\lambda+B)^2 N^2 \delta^2}{N^2 (\min_{j \in \mathcal{J}} \gamma_j)^2}$. This completes the proof of relative compactness.

The proof of the second part follows if the mapping \mathbf{h} is shown to be Lipschitz continuous, i.e., there exists $K > 0$ such that $\rho(\mathbf{h}(\mathbf{u}), \mathbf{h}(\mathbf{v})) \leq K\rho(\mathbf{u}, \mathbf{v})$. For any $\mathbf{u}, \mathbf{v} \in \mathcal{U}$ let the L_1 -distance between \mathbf{u} and \mathbf{v} be defined as $\|\mathbf{u} - \mathbf{v}\| = \sum_{j \in \mathcal{J}} \sum_{\underline{n} \in \mathcal{S}_j} |u_{\underline{n},j} - v_{\underline{n},j}|$. Clearly, we have $\frac{\|\mathbf{u} - \mathbf{v}\|}{(B+1)S} \leq \rho(\mathbf{u}, \mathbf{v}) \leq \|\mathbf{u} - \mathbf{v}\|$, where $S = \sum_{j \in \mathcal{J}} |\mathcal{S}_j|$ and B is as defined before. Thus, these two metrics are equivalent. It is therefore, sufficient to show $\|\mathbf{h}(\mathbf{u}) - \mathbf{h}(\mathbf{v})\| \leq K\|\mathbf{u} - \mathbf{v}\|$ for some $K > 0$. But this is satisfied with $K = 2B + 2\lambda d + 8\lambda S(d-1)$ as can be seen using (12) and the L_1 -norm.

The proof of the third part essentially the same as the proof of Theorem 2 of [16]. We omit the details and mention only the key observation used in the proof which is the fact that as $N \rightarrow \infty$ we have $\mathbf{A}^{(N)} f(\mathbf{u}) \rightarrow \frac{d}{dt} f(\mathbf{x}(t, \mathbf{u}))|_{t=0}$ uniformly in \mathbf{u} for all $f : \mathcal{U} \rightarrow \mathbb{R}$ such that $f(\mathbf{u})$ has bounded partial derivatives of first and second order with respect to each component of \mathbf{u} .

D. Proof of Proposition 2

Consider a *tagged* server of type j and the class- l arrivals that have the tagged server as one of its potential destinations. These arrivals constitute the *potential arrival process* at the tagged server. The probability that the tagged server is sampled at the arrival instant of a job is $\frac{\binom{N-1}{d-1}}{\binom{N}{d}} = \frac{d}{N}$. Thus, due to Poisson thinning, the potential arrival process of class- l jobs to the tagged server is a Poisson process with rate $\frac{d}{N} \times N\lambda_l = d\lambda_l$.

Next, we consider the arrivals that actually join the tagged server. These arrivals constitute the actual arrival process at the server. For finite N , this process is not Poisson since a potential arrival to the tagged server actually joins the server depending on the number of jobs present at the other possible destination server. However, as $N \rightarrow \infty$, due to the asymptotic independence property shown in Proposition 5 the occupancies of the sampled servers become independent of each other. As a result, in equilibrium the actual arrival process converges to a state dependent Poisson process as $N \rightarrow \infty$. Now the arrival rates of the Poisson process, as given in (48), can be computed following the similar line of arguments as in the proof of Lemma 1.