Mean Field Interactions in Heterogeneous Networks

by

Arpan Mukhopadhyay

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Electrical and Computer Engineering

Waterloo, Ontario, Canada, 2015

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In the context of complex networks, we often encounter systems in which the constituent entities randomly interact with each other as they evolve with time. Such random interactions can be described by Markov processes, constructed on suitable state spaces. For many practical systems (e.g. server farms, cloud data centers, social networks), the Markov processes, describing the time-evolution of their constituent entities, become analytically intractable as a result of the complex interdependence among the interacting entities. However, if the 'strength' of these interactions converges to a constant as the size of the system is increased, then in the large system limit the underlying Markov process converges to a deterministic process, known as the *mean field limit* of the corresponding system. Thus, the mean field limit provides a deterministic approximation of the randomly evolving system. Such approximations are accurate for large system sizes.

Most prior works on mean field techniques have analyzed systems in which the constituent entities are identical or *homogeneous*. In this dissertation, we use mean field techniques to analyze large complex systems composed of *heterogeneous* entities.

First, we consider a class of large multi-server systems, that arise in the context of web-server farms and cloud data centers. In such systems, servers with heterogeneous capacities work in parallel to process incoming jobs or requests. We study schemes to assign the incoming jobs to the servers with the goal of achieving optimal performance in terms of certain metrics of interest while requiring the state information of only a small number of servers in the system. To this end, we consider *randomized dynamic job assignment schemes* which sample a small random subset of servers at every job arrival instant and assign the incoming job to one of the sampled servers based on their instantaneous states. We show that for heterogeneous systems, naive sampling of the servers may result in an 'unstable' system. We propose schemes which maintain stability by suitably sampling the servers. The performances of these schemes are studied via the corresponding mean field limits, that are shown to exist. The existence and uniqueness of an asymptotically stable equilibrium point of the mean field is established in every case. Furthermore, it is shown that, in the large system limit, the servers become independent of each other and the stationary distribution of occupancy of each server can be obtained from the unique equilibrium point of the mean field. The stationary tail distribution of server occupancies is shown to have a fast decay rate which suggests significantly improved performance for the appropriate metrics relevant to the application. Numerical studies are presented which show that the proposed randomized dynamic schemes significantly outperform *randomized static schemes* where job assignments are made independently of the server states. In

certain scenarios, the randomized dynamic schemes are observed to be nearly optimal in terms of their performances.

Next, using mean field techniques, we study a different class of models that arise in the context of social networks. More specifically, we study the impact of social interactions on the dynamics of opinion formation in a social network consisting of a large number of interacting social agents. The agents are assumed to be mobile and hence do not have any fixed set of neighbors. Opinion of each agent is treated as a binary random variable, taking values in the set $\{0, 1\}$. This represents scenarios, where the agents have to choose from two available options. The interactions between the agents are modeled using 1) the 'voter' rule and 2) the 'majority' based rule. Under each rule, we consider two scenarios, (1) where the agents are biased towards a specific opinion and (2) where the agents have different propensities to change their past opinions. For each of these scenarios, we characterize the equilibrium distribution of opinions in the network and the convergence rate to the equilibrium by analyzing the corresponding mean field limit. Our results show that the presence of biased agents can significantly reduce the rate of convergence to the equilibrium. It is also observed that, under the dynamics of the majority rule, the presence of 'stubborn' agents (those who do not update their opinions) may result in a *metastable* network, where the opinion distribution of the non-stubborn agents fluctuates among multiple stable configurations.

Acknowledgements

This dissertation is the culmination of four years of research work which would not have been possible without the support of many people.

First and most importantly, I would like to express my sincere gratitude to my supervisor Prof. Ravi R. Mazumdar for his constant support and invaluable advise throughout these years. His positive attitude towards research encouraged me to explore new challenging areas and helped me grow my confidence as a researcher. I shall always cherish my experience of working with a person as knowledgeable and resourceful as him.

I am deeply indebted to Dr. Fabrice Guillemin of Orange Labs, France, with whom I collaborated on the topic discussed in Chapter 3. His thorough reading of my papers and valuable inputs have helped me clear many doubts and correct many subtle mistakes.

I am equally grateful to Prof. Rahul Roy of Indian Statistical Institute, Delhi who suggested me to work on the models dealt in Chapter 4. The discussions I had with him greatly broadened the horizons of my knowledge.

Parts of Chapter 2 and Chapter 3 are the outcome of joint work with A. Karthik who was a post doctoral researcher in our group from Fall 2014 to Spring 2015. I thoroughly enjoyed the discussions I had with him during this period and would like to thank him for the knowledge he shared with me.

I would also like to thank Prof. Patrick Mitran, Prof. Andrew Heunis, Prof. Raouf Boutaba, and Prof. Devavrat Shah for serving on my PhD advisory committee and carefully reading my dissertation. Their comments and suggestions have helped me improve the quality of the presentation of this dissertation.

I am thankful to all my friends and colleagues in Waterloo, whose pleasant company made an otherwise esoteric journey of four years an enjoyable and memorable one.

Finally, I would like to thank my parents and my fiancée, whose constant love and support have sustained me throughout these years. They were the greatest source of strength and inspiration during difficult times. I owe all my accomplishments to them.

Arpan Mukhopadhyay
Waterloo, 2016.

Dedication

This dissertation is lovingly dedicated to my parents and my fiancée.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

A network can be viewed as a collection of interacting entities. The interactions are governed either by a central entity or by a distributed algorithm/protocol running at each constituent entity. Examples of centrally controlled networks include web server farms and cloud data centers. In such systems, the servers interact via a central job dispatcher that assigns incoming jobs to the servers by comparing their loads. In contrast, decentralized network control can be observed in wireless local area networks (WLAN's), wireless sensor networks, and social networks, where each node runs a distributed algorithm that causes the nodes to interact with each other.

In many systems, the interactions among the constituent entities/nodes are random and can be captured through a Markov process that describes the time evolution of the network in its entirety. Characterization of the network's behavior requires analysis of this underlying Markov process. For many systems, an exact analysis of the underlying Markov process becomes intractable due to the complex interdependence among the interacting entities of the system. However, if the 'strength' of the interaction between an individual entity and the rest of the system is 'weak', then the underlying Markov process converges asymptotically to a deterministic process as the number of entities in the system grows to infinity. This limiting deterministic process is known as the *mean field limit* or simply the *mean field* of the system and provides a simpler way to characterize the system's behavior when the number of entities is large. In this dissertation, we analyze large systems of 'weakly' interacting entities through their corresponding mean field limits. In particular, we focus on *heterogeneous* systems in which the statistical properties of the constituent entities are different from each other.

First, we consider a model of multi-server systems (such as shown in Figure 1.1) arising

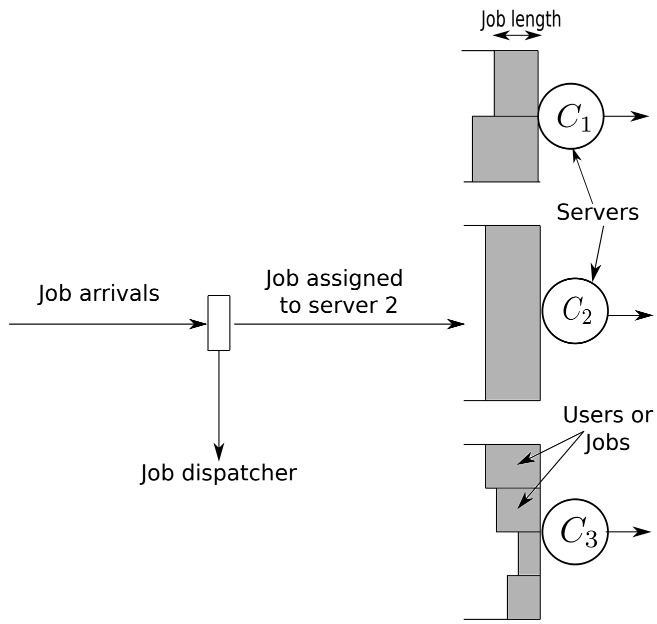

Figure 1.1: Server farm model consisting of $N = 3$ parallel servers. Each job is assigned to the server having the smallest number of unfinished jobs.

in the context of large web server farms and cloud data centers. In such systems, a large number of servers work in parallel to process incoming job requests. Upon arrival, each job request is dispatched by a central job dispatcher to a server for processing. The objective is to assign the incoming jobs to the servers in such a way that certain performance metrics are optimized. Job assignment schemes, which compare the states of some randomly sampled servers at the arrival instant of each job in order to assign the job, effectively cause the servers in the system to interact with each other. We analyze such schemes by analyzing the resulting interactions among the servers using mean field techniques.

Next, we focus on systems (such as shown in Figure 1.2) in which the interactions among the constituent entities occur without any central co-ordination. In particular, we study random interactions between agents in a social network and the effect of such interactions on the diffusion of opinions in the network. The interactions are defined by simple rules in which randomly sampled agents interact with each other to exchange information about their opinions. The global opinion structure that emerges as a result of such interactions is studied using mean field techniques.

The time evolution of both the systems discussed above can be described by suitably constructed Markov processes. Exact characterizations of these processes are extremely

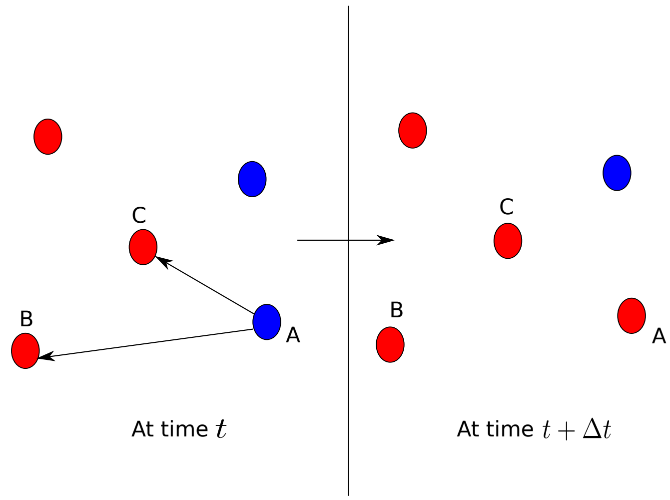

Figure 1.2: Opinion dynamics under the 'majority rule': Different colors represent different opinions. Agent A samples agents B and C at an update instant. It then changes its opinion according to the 'majority rule'.

difficult for finite system sizes. However, it is observed that the 'strength' of the interaction between an individual entity and the rest of the system, as measured by the conditional drift of the individual entity given the entire state of the system, is 'weak' in the sense that it converges to a constant as the number of entities in the system increases. As a result, the underlying Markov processes for these systems, converge in the large system limit to their corresponding mean field limits, that are given by solutions of ordinary differential equations.

Mean field limit of a system serves as a good approximation of the system behavior when the number of entities in the system is large. For example, the steady state behavior of the finite system can be shown to concentrate near the equilibrium points (points in the state space where the time derivative is zero) of the mean field as the number of entities in the system increases. Furthermore, in the large system limit the individual entities of the system become independent of each other. This property is known as the *asymptotic independence* property that enables analysis of the limiting system through the fixed points of certain continuous maps.

Typically, mean field analysis, is carried out for *homogeneous* systems whose constituent entities are identical. But for *heterogeneous* systems - the focus of this dissertation - the analysis becomes more challenging due to the lack of closed form analytical expressions for the equilibrium points of the corresponding mean field limits. Moreover, for such systems, the standard assumptions underlying mean field analysis, such as exchangeability of states among the constituent entities, do not hold. To study these systems, we therefore extend the standard results of mean field theory from the homogeneous case to the heterogeneous case.

1.1 Motivation

Mean field approximations often provide useful insights into the behavior of large systems of 'weakly' interacting entities when an exact analysis of the finite system is intractable. In this dissertation, we use such approximations in two different contexts: (1) to study the performance of randomized job assignment schemes for large web server farms and cloud data centers; (2) to study opinion dynamics in large social networks. Our motivation to analyze these systems stems from their widespread use in practice and the lack of analytical results for the scenarios where the constituent entities are heterogeneous in nature.

1.1.1 Randomized job assignment in heterogeneous multi-server systems

Dynamic content web services (e.g., web search, e-commerce) and cloud computing applications (e.g., Amazon EC2, Microsoft Azure) have seen a rapid growth in the recent years due to their ability to provide scalable and cost effective computing solutions. The prevalence of such applications has triggered the development of large scale web server farms and data centers which can accommodate thousands of servers to meet the increasing demands of user requests. A central problem in such systems is to decide which server an incoming request should be assigned to in order to obtain optimum performance in terms of certain metrics of interest.

In web server farms, the incoming jobs have elastic resource requirements and are processed at the servers in a round-robin fashion with small granularity. This can be well approximated by the processor sharing service discipline [1] where the processing speed of a server is equally shared among the jobs in progress. In this scenario, the objective is to minimize the mean response time of jobs in the system. On the other hand, for cloud data centers providing Infrastructure-as-a-Service (IaaS), the incoming jobs have specific resource requirements and hence must be processed at servers where the required resource is available. Therefore, in this case, the objective is to minimize the probability with which jobs are denied access to requested resource or are *blocked*.

Ideally, a job assignment scheme, which compares the states of all servers in the system to assign every new arrival, should achieve the best performance. However, in large multi-server systems obtaining state information of all the servers at each arrival instant incurs high communication overhead between the servers and the job dispatcher. In such scenarios, the use of *randomized dynamic schemes* which sample only a small random subset of servers at each arrival instant and assign the arrival based on the states of the sampled servers, can drastically reduce the communication overhead. Such schemes are known to perform significantly better than *randomized static schemes* in which the incoming jobs are randomly assigned to the servers without the comparing the server states.

The existing literature on randomized dynamic job assignment schemes [2, 3, 4] assumes that the servers in the system are identical in terms of their capacities. In reality, however, there are servers of varied capacities in web server farms and cloud data centers. In such cases, uniform sampling of servers may not result in the same gains as in the homogeneous case. In this dissertation, our objective is to study randomized dynamic schemes when the servers have heterogeneous capacities. Through our analysis, we conclude that by appropriate sampling of the servers, gains similar to that in the homogeneous case can be obtained even in the heterogeneous case. Our observations are mentioned briefly in Section 1.2 and are discussed in detail in the subsequent chapters of this dissertation.

1.1.2 Opinion dynamics in large social networks with 'biased' and 'stubborn' agents

We are at present in an era where social networks are shaping the opinions of large groups of individuals. Understanding how individual opinions are affected by social interactions and what global opinion structure emerges from such interactions are important in order to make predictions in economics and politics. One of the key challenges in this area is to model social interactions in a way that can lead to meaningful predictions.

Various models have been proposed and analyzed in the literature in this context. One of the models extensively studied is the *voter model* or the *voter rule* [5, 6, 7], where an agent contacts a randomly sampled neighbor at an instant when it decides to update its opinion. The updating agent then copies the opinion of the sampled neighbor. This simple model captures the tendency of individuals to mimic their neighbors and explains why societies often converge to a *consensus*, where all the agents adopt the same opinion. Another common model for social interaction is defined by the *majority rule* [8, 9], where an updating agent samples multiple neighbors and adopts the opinion of the majority of the sampled neighbors. This rule is based on the tendency of individuals to conform with the majority opinion. Under this rule, a fully connected network of agents reaches consensus at a rate faster than that under the voter rule.

In most prior works on voter models and majority rule models, it is implicitly assumed that the agents are 'unbiased' in the sense that they do not have any preference for any of the available choices. It is also assumed that the agents are homogeneous in terms of their propensities to change their past opinions. However, in reality, the agents are often 'biased' towards some specific opinion and/or have different propensities to change their past opinions (i.e., some agents may be more 'stubborn' than others). In such cases, the extension of the existing results is not direct.

In the second part of the dissertation, we address this issue. In particular, we analyze the dynamics of opinion formation under the voter rule and the majority rule assuming two scenarios which represent the cases (1) where the agents are 'biased' towards a specific opinion and (2) where the agents have different propensities to change opinions. We assume that each agent in the network has two available choices represented by the numbers $\{0\}$ and $\{1\}$. Using mean field techniques, we study the equilibrium distribution of opinions in the network as a function of initial opinions of the agents and the number of agents in the network. We also investigate the time required for the network to reach an equilibrium state under the scenarios discussed above.

1.2 Contributions and Outline

This dissertation is organized into two parts. In the first part, we study randomized dynamic job assignment schemes for large, heterogeneous multi-server systems. We start in Chapter 2, with an investigation of randomized schemes to assign arriving requests to servers in a system of parallel processor sharing servers with heterogeneous capacities. The servers are assumed to be grouped into different *types* according to their capacities. The objective is to reduce the mean response time of jobs while requiring the state information of a small number of servers at each job arrival instant. This is achieved by using job assignment schemes which sample a small random subset of servers at every job arrival instant and assign the incoming job to the sampled server with the least number of unfinished jobs. We observe that for heterogeneous systems, the method of sampling the servers plays a key role in determining the performance of such schemes. More specifically, we make the following contributions

- It is shown that uniformly sampling servers from the entire system may drive the system into instability[1] even for arrival rates below the aggregate service rate of the system.

- To recover stability, we propose a hybrid scheme which combines biased sampling across different server types with uniform sampling within the same server type. It is shown that, under the hybrid scheme, the system is stable for all arrival rates below the aggregate service rate of the system, i.e, the hybrid scheme achieves the *maximal stability region*. However, it is observed that the scheme requires the knowledge of the system parameters to achieve the maximal stability region and is therefore not robust to system failures.

- A type-based scheme is proposed where a small number of servers are sampled from each server type at every job arrival instant. The incoming job is then assigned to the sampled server having the least number of unfinished jobs. It is shown that this type-based scheme achieves the maximal stability region for all system parameters and hence is more robust to system failures than the hybrid scheme.

- The performance of the type-based scheme is characterized in the large system limit using mean field techniques. In particular, it is established that the equilibrium point of the mean field is unique and globally attractive. The stationary occupancy distribution of the servers in the finite system is shown to concentrate on this equilibrium point in the large system limit.

- It is shown that, in the limiting system, any finite set of servers behave independently of each other. This is formally known as the *asymptotic independence* property or *propagation of chaos* property. To prove this property in the heterogeneous case, we extend the classical notion of exchangeable random variables to a more general notion of *intra-type exchangeable* random variables.

In Chapter 3, randomized dynamic job assignment schemes are studied for a system of parallel heterogeneous servers each of which holds a finite amount of a resource. The incoming jobs are assumed to have specific resource requirements. Therefore a job can be served at a server only if the required resource is available at the server. In this case, the objective is to design a job assignment scheme that reduces the average blocking probability of jobs while requiring the state information of only a small number of servers at each job arrival instant. To this end, we make the following contributions:

[1] Instability refers to the situation where the mean number of unfinished jobs in the system becomes unbouned.

- We propose a scheme in which an arriving job is assigned to the server having the maximum available resource among a set of randomly sampled servers in the system.
- The performance of the proposed scheme is characterized using mean field techniques. It is shown that the mean field has a unique, asymptotically stable equilibrium point which characterizes the stationary occupancy distribution of the servers in the limiting system.
- Asymptotic independence among the servers is established.
- The tail occupancy distribution of each server in the limiting system is shown to have a fast decay rate which suggests improved performance in terms of reducing blocking.
- Numerical experiments are conducted to show that the proposed scheme results in near minimal average blocking probability under heavy load conditions.

In the second part of the dissertation (Chapter 4), we consider opinion dynamics models for large social networks of interacting agents. The agents are assumed to be mobile and hence do not have any fixed set of neighbors. The opinion of each agent is assumed to be a binary variable taking values in the set $\{0, 1\}$. Each agent is assumed to update its opinion at random instants by interacting with some randomly sampled neighbors. We consider opinion dynamics under the voter rule and the majority rule. Under each rule, we consider two different scenarios. In the first scenario, the agents are assumed to be 'biased' towards a specific opinion. The second scenario assumes that different agents have different propensities to change their opinions. The following are the key observations:

- For the voter model with 'biased' agents we observe that the network reaches a consensus state in a time that is logarithmic in the network size. We also show that the probability with which consensus is achieved on the preferred opinion increases rapidly to 1 as the network size increases.
- For the voter model with agents having different propensities to change opinions, we derive a closed form expression for the probability with which the network reaches consensus on a particular opinion. We also derive an approximation of the mean time to reach consensus.
- In the case of the majority rule model with 'biased' agents, we show that the network converges to a consensus state on the preferred opinion only when the initial fraction of agents having the preferred opinion is more than a certain threshold. The threshold is computed using the mean field limit.

- Finally, for the majority rule model in the presence of 'stubborn' agents (agents who do not update their opinions), we observe that the network may exhibit *metastability* where it fluctuates between multiple stable configurations, spending long intervals in each configuration.

The dissertation is concluded in Chapter 5, where we summarize our work and present some future extensions.

1.3 Mean field techniques: A brief overview

Let us consider a system consisting of N identical interacting particles each of which has a finite state space $\mathcal{S}$. Assume that the time evolution of the entire system of particles can be described by a pure jump Markov process $x^{(N)}(\cdot) = (x_i^{(N)}(\cdot), i \in \mathcal{S})$, where for each $i \in \mathcal{S}$, $x_i^{(N)}(t)$ denotes the fraction of particles in state i at time $t \geq 0$. The process $x^{(N)}(\cdot)$ is called the empirical measure process whose state space is the set of probability measures $\mathcal{P}(\mathcal{S})$ defined on $\mathcal{S}$. We assume that the process $x^{(N)}(\cdot)$ has a unique stationary distribution for each N.

Consider an instant $t \geq 0$ when the system is in state $x^{(N)}(t) = \xi = (\xi_i, i \in \mathcal{S})$. Let $r_{i,j}(\xi)$ denote the rate at which each of the $N\xi_i$ particles in the state $i \in \mathcal{S}$ transits to the state $j \in \mathcal{S} \setminus \{i\}$. The system's state changes from ξ to $\xi + e_j/N - e_i/N$ (where e_k denotes the k^{th} unit vector) when one such transition occurs. Therefore, the total rate at which the process $x^{(N)}(\cdot)$ transits from the state ξ to the state $\xi + e_j/N - e_i/N$ is $N\xi_i r_{i,j}(\xi)$. Clearly, the jump sizes of the process $x^{(N)}(\cdot)$ is $O(1/N)$, whereas the rates at which the jumps occur is $O(N)$.

Convergence to the mean field: Kurtz showed, in [10], that due to the above scaling, the process $x^{(N)}(\cdot)$ converges to a deterministic process $x(\cdot)$ as $N \to \infty$. More specifically, if the initial state $x^{(N)}(0)$ converges in distribution (weak convergence) to a constant ν as $N \to \infty$ and if $r_{i,j} : \mathcal{P}(\mathcal{S}) \to \mathbb{R}$ is Lipschitz continuous for all $i, j \in \mathcal{S}$, then the process $x^{(N)}(\cdot)$ converges in distribution (hence in probability) to a deterministic process $x(\cdot)$ as $N \to \infty$, where $x(\cdot)$ is the unique solution of the following system of ordinary differential equations

$$\dot{x}(t) = x(t)R(x(t)), \tag{1.1}$$

$$x(0) = \nu. \tag{1.2}$$

In the above, $R(\cdot) = [r_{i,j}(\cdot)]_{i,j\in\mathcal{S}}$ denotes the rate matrix of the process $x^{(N)}(\cdot)$. The process $x(\cdot)$ is called the mean field limit of the system. Thus, the above result shows that the path of the stochastic process $x^{(N)}(\cdot)$ tends to concentrate on the path of the deterministic process $x(\cdot)$ as $N \to \infty$. The process $x(\cdot)$ can therefore be used as an approximation of the process $x^{(N)}(\cdot)$ for large values of N.

Convergence of the stationary distribution: An equilibrium point $\pi \in \mathcal{P}(\mathcal{S})$ of the process $x(\cdot)$ is a point which satisfies $\pi R(\pi) = 0$. An equilibrium point π is called a *globally attractive* if $x(t) \to \pi$ as $t \to \infty$ for all $x(0) \in \mathcal{P}(\mathcal{S})$. It can be shown that if the process $x(\cdot)$ has a unique, globally attractive equilibrium point π, then the stationary distribution of the process $x^{(N)}(\cdot)$ (which exists and is unique by assumption) converges in to δ_π as $N \to \infty$, where δ_π denotes the Dirac measure concentrated at the equilibrium point π. Hence, for large N, the steady state of the process $x^{(N)}(\cdot)$ can be approximated by π.

Asymptotic independence: Now we focus on the individual particles in the system. Let $X_n^{(N)}(t) \in \mathcal{S}$ denote the state of the n^{th} particle at time t. Let $X^{(N)}(t) = (X_1^{(N)}(t), X_2^{(N)}(t), \ldots, X_N^{(N)}(t))$. We call $X^{(N)}(t)$ *exchangeable* if the distribution of $X^{(N)}(t)$ remains invariant under the permutation of indices $1 \leq n \leq N$ of the individual particles. We note that since the particles in the system are identical and the system is Markovian, the exchangeability of $X^{(N)}(0)$ implies the exchangeability of $X^{(N)}(t)$ for all $t \geq 0$. It has been shown in [11] that if $X^{(N)}(0)$ in exchangeable and $x^{(N)}(0)$ converges weakly to a deterministic limit as $N \to \infty$, then for any finite set of tagged particles $n_1, n_2, \ldots, n_k$

$$Law(X_{n_1}^{(N)}(\infty), X_{n_2}^{(N)}(\infty), \ldots, X_{n_k}^{(N)}(\infty)) \to \pi^{\otimes k} \text{ as } N \to \infty \tag{1.3}$$

where π is the unique, globally attractive equilibrium point of the mean field $x(\cdot)$ and $X_k^{(N)}(\infty)$ denotes the steady state of the k^{th} particle in the system. Thus, the above result implies that any finite set of particles become independent as $N \to \infty$ and the stationary distribution of each individual particle in the limiting system is given by the unique equilibrium point of the mean field. This property is formally known as the *asymptotic independence* property or the *propagation of chaos* property.

The results discussed above are the key components of mean field analysis. They show how the behavior of a system of weakly interacting particles can be approximated by a deterministic process as the number of particles in the system becomes large. The results also imply that the steady state of the system can be approximated by the equilibrium points of the mean field when the number of particles in the system is large. Such approximations are especially is useful when an exact computation of the stationary distribution of the finite system is intractable. The technique was first introduced in physics [12] to study the evolution of molecules of a dilute gas. The convergence of a sequence of Markov processes to a deterministic process, given by the solution of a system of ordinary differential equations, was formally studied in [10]. The mean field approach soon became a popular methodology in the area of communication networks. In [13], it was used to study the stationary distribution of a closed queuing system in the asymptotic limit as both the number of customers and the number of servers grow to infinity. In [14], it was used to study dynamic alternate routing in large circuit switched networks. In [15], mean sojourn time of jobs was computed in a large network of interacting queues using mean field approximations. More recently, mean field techniques have been used in a variety of contexts such as HTTP flows [16], bandwidth sharing between streaming and file transfers [17], randomized job assignment techniques [2].

We note that the results discussed in this section and in the references mentioned above crucially rely on the fact that the particles in the system are identical and therefore permuting their states does not change their joint distribution. However, in this dissertation our emphasis is on systems in which the constituent particles are of different types. In such cases, permutation of states may affect their joint distribution. This makes the analysis of such systems more complex. We need more general results from the theory of weak convergence of Markov processes to analyze such systems. These results are discussed in Appendix C.

Part I

Randomized Job Assignment in Heterogeneous Multi-Server Systems

Chapter 2

Randomized Job Assignment in Heterogeneous Processor Sharing Systems

In this chapter, we consider the problem of job assignment in a system consisting of a large number of parallel processor sharing servers, categorized into different types according to their capacities. Our objective is to design a job assignment scheme which reduces the mean response time jobs in the system while requiring the state information of only a small subset of servers at each job arrival instant. To this end, we investigate the stability and performance of randomized dynamic job assignment schemes in which a small random subset of servers is sampled at every job arrival instant and the incoming job is assigned to one of the sampled servers based on the states of the sampled servers. We show that in the heterogeneous case, the method of sampling the servers plays an important role in determining the performance of such schemes.

2.1 Introduction

The prevalence of dynamic content web services such as online search, social networking, e-commerce has triggered the growth of large scale web server farms which can accommodate thousands of front-end servers to meet the increasing demand of user requests. Since such systems provide web services that are highly delay sensitive, a small increase in the average response time of jobs may cause a significant loss of revenue and users [18]. Hence, one of the key design issues in such systems is to decide which server an incoming job will be assigned to in order to obtain minimum average response time of jobs. The problem becomes more challenging due to the large size of web server farms which makes the use traditional job assignment schemes inefficient.

Traditionally, in small web server farms, a hardware job dispatcher such as F5 Application Delivery Controller [19, 20] is used to dispatch incoming jobs to the servers. It uses the Join-the-Shortest-Queue (JSQ) scheme that assigns each new arrival to the server having the least number of unfinished jobs in the system. Implementing this scheme requires the knowledge of the states of all the servers in the system at each arrival instant of a new job. This is possible either by probing all the servers at each arrival instant of a new job or by continuously monitoring all the servers in the system at all times. For a large web server farm, where thousands of servers run in parallel to process the incoming jobs, probing all the servers at each job arrival instant introduces a significant delay (due to the large communication overhead between the job dispatcher and the servers) in the routing of the incoming jobs. Continuous monitoring also is not desirable as it is wasteful in terms of energy (since turning off servers of low utilization will require frequent reconfiguration of the job dispatcher). Therefore, for large web server farms, a job assignment scheme, which requires the state information of only a small number of servers at each job arrival instant and yet yields low average response time of jobs, is more desirable.

In *randomized static schemes*, arrivals are assigned to the servers with fixed probabilities, independently of the server states. Thus, randomized static schemes do not require the state information of any server in the system. However, the mean response time of jobs can be further reduced significantly by using *randomized dynamic schemes*, where a small subset of servers is randomly sampled at every job arrival instant and the incoming job is assigned to one of the sampled servers by comparing the states of the sampled servers. Such randomized dynamic schemes, therefore, significantly reduce the mean response time of jobs without increasing the communication overhead between the job dispatcher and the servers significantly.

The the power-of-d scheme or the SQ(d) scheme is an example of randomized dynamic schemes. In the SQ(d) scheme, every arrival is assigned to the server having the least number of unfinished jobs among $d \geq 2$ servers, sampled uniformly at random from the entire system at the arrival instant of the job. It was shown in [2, 3, 21] that for a system of identical first-come-first-serve (FCFS) servers, the SQ(d) scheme results in an exponential reduction in the mean response time of jobs as compared to the randomized static scheme.

However, the above results do not apply to systems consisting of heterogeneous servers, which are more realistic models of web server farms . In heterogeneous systems, a uniformly sampled set of d servers may not contain servers of all types (or capacities). Hence, for such systems, uniform sampling may not yield the same gains as for homogeneous systems. Therefore, an approach based on biased sampling of servers has to be adopted for heterogeneous systems which we explore in this chapter.

2.1.1 Related literature

The study of job assignment schemes for multi-server systems has a long history that dates back to the 1960's. The JSQ scheme was first analyzed in [22, 23] for a system of two parallel FCFS servers assuming Poisson arrivals and exponential service time distribution. The stationary distribution of queue length at each server was found. Optimality of the JSQ scheme, in terms of response time of jobs, was considered in [24, 25]. It was shown that, under FCFS service discipline and service time distributions having decreasing hazard rates[1], JSQ maximizes the number of jobs that depart from the system in a given amount of time.

The study of the JSQ scheme was limited to the FCFS service discipline for a long time until recently Gupta *et al* [26] analyzed the scheme for the processor sharing service discipline, which closely approximates the round robin scheduling of jobs actually employed in web servers [1]. An approximate analysis of the JSQ scheme was presented, assuming a finite number of servers in the system and general job length distributions. It was shown that the JSQ scheme is nearly optimal in terms of minimizing the mean delay in such systems. It was also observed that under the JSQ scheme the stationary distribution of server occupancies is nearly *insensitive* to the type of job length distribution so long as the mean of the distribution remains unchanged.

The concept of randomized dynamic schemes evolved from the seminal work by Azar *et. al* [27] in which the SQ(d) scheme was first proposed in the context of the "balls-and-bins" model. In the balls-and-bins model, n balls are to be sequentially placed in n bins as evenly as possible. The optimal scheme is to check the load on all bins at every step and place the next ball to the least loaded bin (with ties broken arbitrarily). Clearly, this is inefficient when n is large. However, if each ball is placed in the least loaded of d randomly sampled bins (with ties broken uniformly at random), then it was shown that (with high probability) the maximum load in a bin (after completion) is $(1 + o(1)) \ln n / \ln \ln n$ for

[1]For a service distribution, having cumulative distribution function (CDF) $F(\cdot)$, the hazard rate $h(\cdot)$ is given by $h(t) = f(t)/(1 - F(t))$, where $f(\cdot)$ denotes the probability density function (PDF) of the distribution.

$d = 1$, and $\ln \ln n / \ln d + O(1)$ for $d \geq 2$. Thus, an exponential reduction in the maximum load can be obtained by increasing the number of samples from $d = 1$ to $d = 2$.

Motivated by the above observation, a dynamic version of the problem was analyzed in [2, 3]. It was assumed that jobs having exponentially distributed job lengths arrive at a bank of N identical FCFS queues according to a Poisson process with rate $N\lambda$. The system was studied in the limit as $N \to \infty$ using mean field techniques. For the SQ(d) scheme with $d \geq 2$, [2] showed, using the theory of operator semigroups, that the equilibrium tail distribution queue sizes decay super-exponentially in the large system limit. Mitzenmacher in [3, 21] derived the same result using an extension of Kurtz's theorem [28]. *Chaoticity on path space* or *asymptotic independence of the queue length processes* was established in [29] using empirical measures on the path space. Results of [2] were generalized to the case of open Jackson networks in [30].

The case, where queues have different service rates, was considered in [31]. It was shown that, under the SQ(d) scheme, the system is not stable for all arrival rates below the aggregate service rate of the system. The stability region, however, was not characterized explicitly. To recover the stability region, a variant of the SQ(d) scheme based on memory was proposed. In this scheme, after an arrival has been assigned, the least loaded of d sampled servers for that arrival is kept in the memory and is used as one of the d potential destination servers for the next arrival. It was shown that this memory based scheme can support all arrival rates below the aggregate service rate of the system. However, our results show that the memory based scheme do not perform well in comparison to the schemes proposed in this chapter.

Recently, in [32], the SQ(d) scheme was analyzed under more general service disciplines and service time distributions. It was shown that in the case of FCFS discipline and power-law service time distribution, the equilibrium tail distribution of queue sizes decay super-exponentially, exponentially, or just polynomially, depending on the power-law exponent and the number of choices, d. The stability of more general randomized shortest queue based schemes for non-idling service disciplines was analyzed in [33], which derived a sufficient condition under which such networks are stable. Asymptotic independence of servers in equilibrium was conjectured under any local service disciplines and general service time distributions in [34]. The conjecture was proved only for FCFS service discipline and service time distributions having decreasing hazard rate (DHR) functions.

The tradeoff between sampling cost of servers and the expected sojourn time seen by a customer under the SQ(d) scheme was studied in [35] for FCFS queues. A game theoretic framework was proposed. It was shown that for arrival rates within the stability region of the network, a symmetric Nash equilibrium for identical customers exists in which each customer chooses a fixed number of queues to sample.

2.1.2 Contributions

Analysis of randomized dynamic job assignment schemes has been mostly restricted to the homogeneous case, where the servers have identical capacities. However, in reality, a web server farm contains servers of varied capacities for which the above results do not apply. In this chapter, our focus will be on analyzing the performance of randomized dynamic job assignment schemes for heterogeneous processor sharing systems.

We first characterize the stability and performance of the SQ(d) scheme under the heterogeneous scenario. In particular, we show that the stability region (the set of arrival rates for which the underlying Markov process is positive recurrent) for the heterogeneous system operating under the SQ(d) scheme is a subset of the maximal stability region, defined as the set of arrival rates below the aggregate capacity of the system.

To recover the loss in the stability region, we then propose a hybrid scheme which combines probabilistic routing across different server types with the SQ(d) routing within servers of the same type. We show that the proposed hybrid scheme achieves the maximal stability region but only with the knowledge of the system parameters. We also obtain the optimal routing probabilities for which the mean response time of jobs in the system is minimized under this scheme.

Next, we propose a type-based scheme in which a small subset of servers is sampled from every server type at each job arrival instant. This ensures that servers of all types are present in the sampled set of potential destination servers for each arrival. The job is then assigned to one of the sampled servers based on their instantaneous loads. We show that, unlike the hybrid scheme discussed above, this scheme achieves the maximal stability region even without the knowledge of the system parameters. We characterize the performance of this scheme in the large system limit using mean field techniques.

Finally, we study the performance of the above schemes numerically. The proposed schemes are found to outperform other existing randomized schemes. Specifically, the hybrid scheme and the type-based scheme are observed to result in the lowest mean response times of jobs among all the randomized schemes considered. Another important observation, made through the numerical studies, is the fact that in the limiting system, the mean response time of jobs does not depend on the type of job length distribution so long as the mean of the distribution remains unchanged. We refer to this property as the *asymptotic insensitivity* of the system.

2.1.3 Organization

The rest of the chapter is organized as follows. In Section 2.2, we introduce the system model, the job assignments schemes considered throughout the chapter, and the notations used to analyze them. In Section 2.3 we review some of the existing results in detail. In Sections 2.4, 2.5, and 2.6 we present the detailed analysis of the randomized dynamic schemes in the heterogeneous scenario. In section 2.7, we report the results of our numerical studies. Finally, the chapter is concluded in Section 2.8.

2.2 System Model

We consider a system consisting of N parallel processor sharing (PS) servers with heterogeneous service rates or capacities. The capacity, C, of a server is defined as the time-rate at which it processes a single job assigned to it. If $q(t)$ jobs are present at a server of capacity C at time t, then the rate at which each job is processed at time t is given by $C/q(t)$. We assume that the servers are divided into $M(<< N)$ different *types* according to their capacities. Let $\mathcal{J} = \{1, 2, \dots, M\}$ denote the index set of server types. The capacity of each server of type $j \in \mathcal{J}$ is denoted by C_j. Let $\mathcal{C} = \{C_1, C_2, \dots, C_M\}$ denote the set of all server capacities. We assume, without loss of generality, that the server capacities are ordered in the following way

$$C_1 \leq C_2 \leq \dots C_M. \tag{2.1}$$

For each $j \in \mathcal{J}$, the proportion of servers with capacity C_j is assumed to be fixed and is denoted by γ_j $(0 \leq \gamma_j \leq 1)$. Clearly, $\sum_{j=1}^{M} \gamma_j = 1$.

Jobs are assumed to arrive at the system according to a Poisson process with rate $N\lambda$. Each job brings a random amount of work, independent and exponentially distributed with a finite mean $1/\mu$. The inter-arrival times and the job lengths are assumed to be independent of each other. Upon arrival, a job is assigned to one of the N servers where the job stays till the completion of its service after which it leaves the system. We consider the following randomized schemes to assign the incoming jobs to the servers.

2.2.1 Scheme 1: The randomized static scheme

As a baseline, we consider a scheme that assigns the incoming jobs to the servers with fixed probabilities, independent of the current states of the servers. Upon arrival of a job a server type $j \in \mathcal{J}$ is first chosen with probability p_j. The job is then assigned to any one of the $N\gamma_j$ servers of the chosen type with equal probability[2]. The probabilities p_j, $j \in \mathcal{J}$, then must satisfy $\sum_{j\in\mathcal{J}} p_j = 1$. Clearly, in this scheme, no communication is required between the job dispatcher and the servers as the job assignments are made independently of the states of the servers.

2.2.2 Scheme 2: The power-of-d or SQ(d) scheme

In this scheme, a subset of $d \geq 2$ servers is sampled uniformly at random with replacement[3] from the set of N servers at each arrival instant. These d sampled servers are the called the *potential destination* servers for the incoming arrival. The incoming job is assigned to the server having the least number of unfinished jobs among the d potential destination servers. In case of a tie within sampled servers of the same type, the tie is broken by choosing any one of the tied servers uniformly at random. In case of a tie between sampled servers of different types, the tie is broken by choosing the server having the highest capacity. We shall refer to the server selected after tie breaking as the *actual destination* server for the incoming arrival.

2.2.3 Scheme 3: The hybrid SQ(d) scheme

In this scheme, upon arrival of a new job, a server type $j \in \mathcal{J}$ is chosen with probability p_j, where the probabilities p_j, $j \in \mathcal{J}$, satisfy $\sum_{j\in\mathcal{J}} p_j = 1$. The job dispatcher then samples d servers uniformly at random (with replacement) from selected server type. Finally, the incoming job is assigned to the server having the least number of unfinished jobs among the d sampled servers. Ties are broken uniformly at random. Hence, this scheme combines probabilistic routing across different server types with the SQ(d) routing within servers of the same type.

2.2.4 Scheme 4: The type-based scheme

In this scheme, upon arrival of a job, d_j servers of type j are sampled uniformly at random (with replacement) for all $j \in \mathcal{J}$. The job is then assigned to the server having the least

[2]Note that servers of the same type are chosen with the same probability since, by symmetry, choosing servers of the same type with different probabilities cannot yield lower mean response time of jobs.

[3]We have seen that sampling with or without replacement yield the same results in the limit as $N \to \infty$

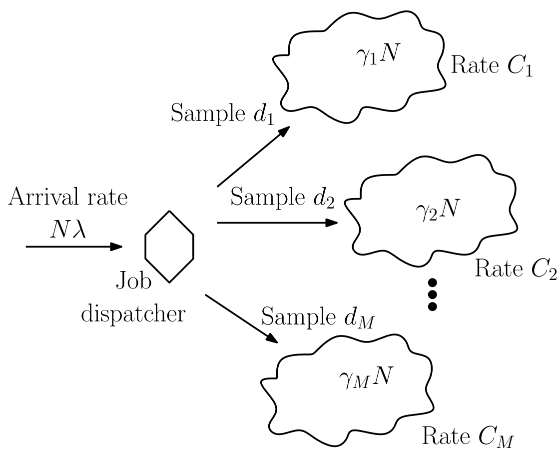

Figure 2.1: The type-based scheme: System consisting of N parallel processor sharing (PS) servers, categorized into M types. There are $\gamma_j N$ servers of type j, each of which has a capacity or rate C_j. Arrivals occur according to a Poisson process with rate $N\lambda$. For each arrival, the job dispatcher samples d_j servers of type j and routes the arrival to the least loaded of the sampled servers.

instantaneous occupancy among the $\sum_{j\in\mathcal{J}} d_j$ sampled servers. Ties among servers of the same type are broken uniformly at random and ties across server types are broken by selecting the server type having the highest capacity. We note that, in this scheme, servers of each type are present in the sampled set of potential destination servers. This is unlike the SQ(d) scheme, where all server types may not be present in the sampled set of servers for an arrival. A schematic diagram describing this scheme is given in Figure 2.1.

2.2.5 Additional notations

Throughout the analysis we shall use $\mathbb{Z}$, $\mathbb{Z}_+$, $\mathbb{N}$, $\mathbb{R}$ to denote the set of all integers, the set of non-negative integers, the set of positive integers, and the set of real numbers, respectively. We define the following sequence spaces

$$\bar{\mathcal{U}} = \{(g_n, n \in \mathbb{Z}_+) : g_0 = 1, g_n \geq g_{n+1} \geq 0 \text{ for all } n \in \mathbb{Z}_+\}. \tag{2.2}$$

$$\mathcal{U}_j^{(N)} = \left\{(g_n, n \in \mathbb{Z}_+) \in \bar{\mathcal{U}} : N\gamma_j g_n \in \mathbb{N} \text{ for all } n \in \mathbb{Z}_+\right\}, j \in \mathcal{J} \tag{2.3}$$

$$\mathcal{U} = \{(g_n, n \in \mathbb{Z}_+) \in \bar{\mathcal{U}} : \sum_{n\geq 0} g_n < \infty\}. \tag{2.4}$$

We shall mainly be using the product spaces $\bar{\mathcal{U}}^M$, $\mathcal{U}^{(N)} = \prod_{j\in\mathcal{J}} \mathcal{U}_j^{(N)}$, and $\mathcal{U}^M$ in our analysis. An element $\mathbf{u} = (u_{n,j}, n \in \mathbb{Z}_+, j \in \mathcal{J})$ is said to belong to $\bar{\mathcal{U}}^M$, $\mathcal{U}^{(N)}$ or $\mathcal{U}^M$ if, for each $j \in \mathcal{J}$, the sequence $(u_{n,j}, n \in \mathbb{Z}_+)$ belongs to $\bar{\mathcal{U}}$, $\mathcal{U}_j^{(N)}$, or $\mathcal{U}$, respectively. We define a metric ω on the space $\bar{\mathcal{U}}^M$ as follows

$$\omega(\mathbf{u}, \mathbf{v}) = \sup_{j\in\mathcal{J}} \sup_{n\in\mathbb{Z}_+} \left| \frac{u_{n,j} - v_{n,j}}{n+1} \right|, \text{ for all } \mathbf{u}, \mathbf{v} \in \bar{\mathcal{U}}^M. \tag{2.5}$$

The space $\bar{\mathcal{U}}^M$ is compact (and hence complete and separable) under the metric ω (proof given in Appendix A).

By δ_x we will denote the Dirac measure concentrated at point x. Weak convergence (convergence in distribution) of a sequence of probability measures $(\nu_n)_n$ (sequence of random variables $(X_n)_n$) to a probability measure ν (random variable X) is denoted by $\nu_n \Rightarrow \nu$ ($X_n \Rightarrow X$). For the formal definition of weak convergence the reader is referred to Appendix C. For a measure space $(H, \mathcal{H}, \mu_H)$ and a μ_H-integrable function $f : H \to \mathbb{R}$, we define duality brackets as $\langle f, \mu_H\rangle = \int f d\mu_H$.

2.3 A brief review of previous results

In this section, we briefly review the existing results for the randomized static scheme in the heterogeneous case and the SQ(d) scheme in the homogeneous case.

2.3.1 The randomized static scheme

We now restate the results of [36] in terms of the notations defined in Section 2.2. In the randomized static scheme, a job is assigned to a server with a fixed probability, independent of the instantaneous states of the servers in the system. Hence, under this scheme, the system reduces to a set of independent M/M/1 processor sharing servers working in parallel.

Stability analysis: We first consider the stability of the system operating under the randomized static scheme. The stability region is defined to be the set of arrival rates for which the mean sojourn time of jobs in the system is bounded. Equivalently, it is set of arrival rates for which the underlying Markov process describing the time evolution of the system is *positive recurrent*. Proposition 1 of [36] states that there always exist probabilities p_j, $j \in \mathcal{J}$, for which the system is stable under the randomized static scheme if and only if the following condition holds:

$$\lambda \in \Lambda = \left\{ 0 \leq \lambda < \mu \sum_{j \in \mathcal{J}} \gamma_j C_j \right\}. \tag{2.6}$$

Hence, the stability region of the system operating under the randomized static scheme is given by the set Λ. We refer to the set Λ as the *maximal stability region* of the system because for $\lambda > \mu \sum_{j \in \mathcal{J}} \gamma_j C_j$ the system is unstable under any job assignment scheme (see [33]).

It was also shown in [36] that a choice of the routing probabilities for which the system operating under the randomized static scheme is stable for all $\lambda \in \Lambda$ is given by

$$p_j = \frac{\gamma_j C_j}{\sum_{i \in \mathcal{J}} \gamma_i C_i} \text{ for all } j \in \mathcal{J} \tag{2.7}$$

The optimal routing probabilities: The above choice of the routing probabilities is not optimal in terms of minimizing the mean sojourn time of jobs under the randomized static scheme. The routing probabilities p_j^*, $j \in \mathcal{J}$, for which the mean sojourn time of jobs in the system is minimized under the randomized static scheme were found in Theorem 1 of [36] by solving a convex optimization problem. The optimal routing probabilities were found to depend on the arrival rate λ. In particular, it was found that, depending on the arrival rate of jobs, all server types may not be used in the optimal randomized static scheme. The only server types used in the optimal static scheme was found to be $\mathcal{J}_{\text{opt}} = \{j^*, j^*+1, \ldots, M\} \subseteq \mathcal{J}$, where j^* is given by

$$j^* = \inf \left\{ j \in \mathcal{J} : \frac{1}{\sqrt{C_j}} < \frac{\sum_{i=j}^{M} \gamma_i \sqrt{C_i}}{\sum_{i=j}^{M} \gamma_i C_i - \frac{\lambda}{\mu}} \right\}. \tag{2.8}$$

Furthermore, the optimal loads $\vec{\rho}^* = (\rho_1^*, \rho_2^*, \ldots, \rho_M^*)$ were found to be

$$\rho_i^* = \begin{cases} 1 - \sqrt{\frac{1}{C_i} \frac{\sum_{k=j^*}^{M} \gamma_k C_k - \frac{\lambda}{\mu}}{\sum_{k=j^*}^{M} \gamma_k \sqrt{C_k}}}, & \text{if } i \in \mathcal{J}_{\text{opt}} \\ 0, & \text{otherwise.} \end{cases} \tag{2.9}$$

The optimal routing probability p_j^*, is related to the optimal load ρ_j^* as $\rho_j^* = p_j^* \lambda / \gamma_j \mu C_j$ for each $j \in \mathcal{J}$.

Applying Little's law, the mean sojourn time of jobs in the system under the optimal static scheme can be obtained as

$$\bar{T}^{\text{static}}(\lambda) = \frac{1}{\lambda}\sum_{j=1}^{M}\gamma_j\left(\frac{\rho_j^*}{1-\rho_j^*}\right), \tag{2.10}$$

Remark 1. Specializing the above results to the homogeneous case with $M = 1$, $C_1 = C$, and $\gamma_1 = 1$, we have $\mathcal{J}_{\text{opt}} = \{1\}$ and $\rho^* = \rho_1^* = \lambda/\mu C$. Hence, the mean sojourn time of jobs in the homogeneous system is given by

$$\bar{T}^{\text{static}}_{\text{homo}}(\lambda) = \left(\frac{1}{\mu C - \lambda}\right) = \frac{1}{\lambda}\sum_{n=1}^{\infty}\left(\frac{\lambda}{\mu C}\right)^n \tag{2.11}$$

2.3.2 The SQ(d) scheme in the homogeneous case

An exact analysis of the SQ(d) scheme for any finite value of N is extremely difficult since under the SQ(d) scheme the servers in the system are dependent on each other. However, in the large N limit, the SQ(d) scheme was analyzed using mean field techniques in [3, 2]. We now recall their results. To state these results, we assume $M = 1$, $C_1 = C$, and $\gamma_1 = 1$, which corresponds to the homogeneous system with N identical servers having capacity C. Hence, in this case, the maximal stability region can be obtained from (2.6) as $\Lambda = \{0 \le \lambda < \mu C\}$.

For the homogeneous system of size N, operating under the SQ(d) scheme, let $x_n^{(N)}(t)$ denote the fraction of servers with at least n unfinished jobs at time t. Clearly, $\mathbf{x}^{(N)}(\cdot) = (x_n^{(N)}(\cdot), n \in \mathbb{Z}_+)$ defines a Markov process.

Stability analysis: Using simple coupling arguments, it was shown in [2, 3] that the process $\mathbf{x}^{(N)}(\cdot)$, defined above, is positive recurrent if $\lambda \in \Lambda$. Hence, the SQ(d) scheme achieves the maximal stability region for homogeneous systems.

Mean field analysis: Analyzing the Markov process $\mathbf{x}^{(N)}(\cdot)$ is intractable for finite N due to the dependence among servers in the SQ(d) scheme. However, using mean field analysis it was shown in [2, 3] that if $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0$ as $N \to \infty$ for some constant $\mathbf{u}_0 \in \bar{\mathcal{U}}$, then $\mathbf{x}^{(N)}(\cdot) \Rightarrow \mathbf{x}(\cdot) = \{x_n(\cdot), n \in \mathbb{Z}_+\}$[4] where the process $\mathbf{x}(\cdot)$ is a deterministic process lying in the space $\bar{\mathcal{U}}$ and is given by the unique solution of the following system of differential equations.

[4]In this case $\Rightarrow$ denotes the weak convergence of the process $\mathbf{x}^{(N)}(\cdot)$ to the process $\mathbf{x}(\cdot)$ as $N \to \infty$.

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{2.12}$$

$$\dot{x}_0(t) = 0, \tag{2.13}$$

$$\dot{x}_n(t) = \lambda \left(x_{n-1}^d(t) - x_n^d(t) \right) - \mu C \left(x_n(t) - x_{n+1}(t) \right), n \geq 1. \tag{2.14}$$

Thus, the above result shows that as $N \to \infty$ the process $\mathbf{x}^{(N)}(\cdot)$ tends to concentrate on the path of the deterministic process $\mathbf{x}(\cdot)$. The process $\mathbf{x}(\cdot)$ is called the mean field limit of the system.

An equilibrium point of the mean field is a point $\mathbf{P}$ in the state space of the process $\mathbf{x}(\cdot)$ such that $\dot{\mathbf{x}}(t) = \mathbf{0}$ if $\mathbf{x}(t) = \mathbf{P}$. It is easy to see that the process $\mathbf{x}(\cdot)$, as defined above, has a unique equilibrium point $\mathbf{P} = (P_n, n \in \mathbb{Z}_+)$ in the space $\mathcal{U}$ given by

$$P_n = \left(\frac{\lambda}{\mu C} \right)^{\frac{d^n - 1}{d-1}}. \tag{2.15}$$

It was shown in [2] that for $\lambda \in \Lambda$, $x_n^{(N)}(\infty) \Rightarrow P_n$ as $N \to \infty$, where $x_n^{(N)}(\infty)$ denotes the equilibrium fraction of servers in the system having at least n unfinished jobs. Thus, the above result shows that the steady state $\mathbf{x}^{(N)}(\infty)$ of the process $\mathbf{x}^{(N)}(\cdot)$ tends to concentrate near the equilibrium point $\mathbf{P}$ of the mean field as $N \to \infty$.

Using the above results, the mean response time of jobs in the limiting system can be computed as follows. Since $x_n^{(N)}(\infty) \Rightarrow P_n$ and $0 \leq x_n^{(N)}(\infty) \leq 1$ is bounded, we have $\lim_{N\to\infty} \mathbb{E}[x_n^{(N)}(\infty)] = P_n$. Now, for a given system of size N, the expected number of jobs in the system at equilibrium is $N\mathbb{E}[\sum_{n=1}^{\infty} x_n^{(N)}(\infty)] = N \sum_{n=1}^{\infty} \mathbb{E}[x_n^{(N)}(\infty)]$ (which is finite if $\lambda \in \Lambda$ due to stability). Therefore, by Little's law, the mean sojourn time of jobs in the finite system of size N for $\lambda \in \Lambda$ is given by

$$\bar{T}_{N,\text{homo}}^{\text{sqd}}(\lambda) = \frac{N \sum_{n=1}^{\infty} \mathbb{E}\left[x_n^{(N)}(\infty)\right]}{N\lambda} = \frac{\sum_{n=1}^{\infty} \mathbb{E}\left[x_n^{(N)}(\infty)\right]}{\lambda} \tag{2.16}$$

Taking the limit of the above equation as $N \to \infty$, we obtain the mean sojourn time of jobs in the limiting system to be

$$\bar{T}_{\text{homo}}^{\text{sqd}}(\lambda) = \frac{1}{\lambda} \sum_{n=1}^{\infty} P_n = \frac{1}{\lambda} \sum_{n=1}^{\infty} \left(\frac{\lambda}{\mu C} \right)^{\frac{d^n - 1}{d-1}}. \tag{2.17}$$

Comparing the above expression with the expression of mean sojourn time in (2.11) it is easy to see that for any $d \geq 2$, there is a significant reduction in the mean sojourn time of jobs in the SQ(d) scheme as compared to the that in the randomized static scheme for the homogeneous case. The reduction is due to the fact that P_n decays super-exponentially with the increase in n for $d \geq 2$.

2.4 Analysis of the SQ(d) scheme in the heterogeneous case

We now generalize the results of [2, 3] to the heterogeneous case. To state our main results, we first introduce the following notations.

For the model described in Section 2.2, let $x_{k,j}^{(N)}(t)$, $k \in \mathbb{Z}_+, j \in \mathcal{J}, t \geq 0$ denote the fraction of type j servers having at least k unfinished jobs at time t in the finite system of size N. Define $\mathbf{x}^{(N)}(\cdot) = (x_{k,j}^{(N)}(\cdot), k \in \mathbb{Z}_+, j \in \mathcal{J})$ to be the process describing the time evolution of the system. Clearly, under the assumptions of Poisson arrivals and exponential job length distribution, the process $\mathbf{x}^{(N)}(\cdot)$ is Markov.

Let N^* denote the smallest positive integer (> 2) such that $\gamma_j N^*$ is a positive integer for all $j \in \mathcal{J}$. Now, let Λ_k, $k \in \mathbb{N}$, denote the stability region of the system operating under the SQ(d) scheme when $N = kN^*$. In other words, when there are $N = kN^*$ servers in the system, the process $\mathbf{x}^{(N)}(\cdot)$ is positive recurrent if $\lambda \in \Lambda_k$.

We also define the set Λ_∞ as follows

$$\Lambda_\infty = \left\{ 0 \leq \lambda < \mu \min_{1 \leq m \leq M} \left\{ \frac{\left(\sum_{j=1}^{m} \gamma_j C_j\right)}{\left(\sum_{j=1}^{m} \gamma_j\right)^d} \right\} \right\} \tag{2.18}$$

From the above expression and (2.6), it is easy to see (by putting $m = M$) that $\Lambda_\infty \subseteq \Lambda$.

Main Results: Our main results are the following:

1. For Λ_k, $k \in \mathbb{N}$, as defined above, and Λ as given in (2.6), we have $\Lambda \supseteq \Lambda_1 \supseteq \Lambda_2 \supseteq \ldots$. Furthermore, $\Lambda_\infty \subseteq \cap_{k=1}^{\infty} \Lambda_k$.
2. If $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \bar{\mathcal{U}}^M$ as $N \to \infty$, then $\mathbf{x}^{(N)}(\cdot) \Rightarrow \mathbf{x}(\cdot)$ as $N \to \infty$, where the process $\mathbf{x}(\cdot) = (x_{k,j}(\cdot), k \in \mathbb{Z}_+, j \in \mathcal{J})$, lying in the space $\bar{\mathcal{U}}^M$, is given by the unique solution of the following system of differential equations:

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{2.19}$$

$$\dot{\mathbf{x}}(t) = \mathbf{h}(\mathbf{x}(t)). \tag{2.20}$$

The mapping $\mathbf{h} : \bar{\mathcal{U}}^M \to (\mathbb{R}^{\mathbb{Z}_+})^M$ is given by

$$h_{0,j}(\mathbf{x}) = 0, \text{ for } j \in \mathcal{J}, \tag{2.21}$$

$$h_{n,j}(\mathbf{x}) = \frac{\lambda}{\gamma_j}\left[\left(\sum_{i=1}^{j}\gamma_i x_{n-1,i} + \sum_{i=j+1}^{M}\gamma_i x_{n,i}\right)^d - \left(\sum_{i=1}^{j-1}\gamma_i x_{n-1,i} + \sum_{i=j}^{M}\gamma_i x_{n,i}\right)^d\right]$$
$$- \mu C_j\left(x_{n,j} - x_{n+1,j}\right), \text{ for } n \geq 1, j \in \mathcal{J}. \tag{2.22}$$

3. If $\lambda \in \Lambda_\infty$, then there exists a unique solution, $\mathbf{P}$, of $\mathbf{h}(\mathbf{P}) = \mathbf{0}$ in the space $\mathcal{U}^M$. The point $\mathbf{P} \in \mathcal{U}^M$ is called the equilibrium point of the process $\mathbf{x}(\cdot)$ described by (2.19)-(2.20). Moreover, for all $\mathbf{u}_0 \in \mathcal{U}^M$ we have

$$\lim_{t\to\infty} \mathbf{x}(t, \mathbf{u}_0) = \mathbf{P}, \tag{2.23}$$

where $\mathbf{x}(\cdot, \mathbf{u}_0)$ denotes the process $\mathbf{x}(\cdot)$ started at $\mathbf{x}(0) = \mathbf{u}_0$.

4. If $\lambda \in \Lambda_\infty$, then $\mathbf{x}^{(N)}(\infty) \Rightarrow \mathbf{P}$ as $N \to \infty$, where $\mathbf{x}^{(N)}(\infty) = \lim_{t\to\infty}\mathbf{x}^{(N)}(t)$, denotes the random variable representing the equilibrium state of the finite system of size N.

Remark 2. The first result implies that for any finite N, the stability region of the system operating under the SQ(d) scheme is a subset of the maximal stability region Λ given by (2.6). Thus, in the heterogeneous case, the stability region under the SQ(d) scheme is smaller than that under the randomized static scheme. The reason behind the reduction of the stability region can be intuitively explained as follows.

In the SQ(d) scheme servers are sampled uniformly at random at each arrival instant. Hence, the servers with higher capacities are chosen with the same probability as the servers with lower capacities. In this respect, the randomized static scheme provides more flexibility by allowing to choose servers of different capacities with different probabilities.

Remark 3. The second result implies that if the initial state of the finite system concentrates on some point $\mathbf{u}_0 \in \bar{\mathcal{U}}^M$, then as $N \to \infty$, the process $\mathbf{x}^{(N)}(\cdot)$ concentrates on the path of the deterministic process $\mathbf{x}(\cdot)$ taking values in the space $\bar{\mathcal{U}}^M$. The process $\mathbf{x}(\cdot)$

is called the mean field limit of the heterogeneous system under the SQ(d) scheme. The condition $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \bar{\mathcal{U}}^M$ is satisfied, for example, if the system is initially empty or if for each $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$ the fraction of servers of type j with at least k unfinished jobs in the initial system, denoted by $x_{k,j}^{(N)}(0)$, is a constant independent of N. The weak convergence of the process $\mathbf{x}^{(N)}(\cdot)$ to the deterministic process $\mathbf{x}(\cdot)$ also implies the following convergences:

- For each $t \geq 0$, $\mathbf{x}^{(N)}(t) \Rightarrow \mathbf{x}(t)$ as $N \to \infty$.
- For each $t \geq 0$, $j \in \mathcal{J}$, and $k \in \mathbb{Z}_+$, $x_{k,j}^{(N)}(t) \Rightarrow x_{k,j}(t)$ as $N \to \infty$.
- For each $t \geq 0$, $j \in \mathcal{J}$, and $k \in \mathbb{Z}_+$, $\mathbb{E}[x_{k,j}^{(N)}(t)] \to \mathbb{E}[x_{k,j}(t)]$ as $N \to \infty$.

Remark 4. The third result implies that if $\lambda \in \Lambda_\infty$ then there exists a unique equilibrium point $\mathbf{P}$ of the process $\mathbf{x}(\cdot)$ in the space $\mathcal{U}^M$. It further states that starting from any initial state $\mathbf{u}_0 \in \mathcal{U}^M$, the process $\mathbf{x}(\cdot)$ reaches its unique equilibrium point $\mathbf{P} \in \mathcal{U}^M$ as $t \to \infty$. An equilibrium point $\mathbf{P}$ with the above property is referred to as a *globally asymptotically stable* equilibrium point.

Remark 5. Since $\Lambda_\infty \subseteq \Lambda_k$, for all k, we have that the process $\mathbf{x}^{(N)}(\cdot)$ is positive recurrent for all $\lambda \in \Lambda_\infty$. Equivalently, there exists a unique equilibrium distribution of the process $\mathbf{x}^{(N)}(\cdot)$ for every N if $\lambda \in \Lambda_\infty$. Thus, for $\lambda \in \Lambda_\infty$ the steady state $\mathbf{x}^{(N)}(\infty) = (x_{k,j}^{(N)}(\infty), k \in \mathbb{Z}_+, j \in \mathcal{J})$ is finite and is distributed according to the unique stationary distribution of the process $\mathbf{x}^{(N)}(\cdot)$. The last result then implies that for each $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$, the equilibrium fraction of type j servers having at least k unfinished jobs in the system of size N, converges to $P_{k,j}$ as $N \to \infty$. Thus, the quantity $P_{k,j}$ can be treated as the equilibrium fraction of type j servers having at least k unfinished jobs in the limiting system. Since $\mathbf{P}$ satisfies $\mathbf{h}(\mathbf{P}) = \mathbf{0}$, we have from (2.21) and (2.22) that $P_{0,j} = 1$ for all $j \in \mathcal{J}$ and

$$P_{k+1,j} - P_{k+2,j} = \frac{\lambda}{\mu C_j \gamma_j} \left[\left(\sum_{i=1}^{j} \gamma_i P_{k,i} + \sum_{i=j+1}^{M} \gamma_i P_{k+1,i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k,i} + \sum_{i=j}^{M} \gamma_i P_{k+1,i} \right)^d \right], \tag{2.24}$$

for all $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$. The quantities $P_{k,j}$, $k \geq 0$, $j \in \mathcal{J}$, can be computed by solving the recursive relation given above. These quantities can then be used to compute the mean sojourn time of jobs in the limiting system by using Little's law. Following similar steps as outlined for the homogeneous case, it is easy to see that the mean sojourn time of jobs in the limiting heterogeneous system is given by

$$\bar{T}^{\text{sqd}}(\lambda) = \frac{1}{\lambda}\sum_{j=1}^{M}\gamma_j\sum_{k=1}^{\infty}P_{k,j}. \tag{2.25}$$

We now provide the detailed proof of the first result. The proofs of the other results are similar to those of the corresponding results for the type based scheme, discussed in Section 2.6. We, therefore, do not repeat them here.

Theorem 2.4.1. *For Λ_k, $k \in \mathbb{N}$, as defined above, and Λ as given in (2.6), we have*

$$\Lambda \supseteq \Lambda_1 \supseteq \Lambda_2 \supseteq \ldots \tag{2.26}$$

Furthermore, $\Lambda_\infty \subseteq \cap_{k=1}^{\infty}\Lambda_k \subseteq \bar{\Lambda}_\infty$, where $\bar{\Lambda}_\infty$ denotes the closure of Λ_∞ and is given by

$$\bar{\Lambda}_\infty = \left\{0 \le \lambda \le \mu \min_{1\le m\le M}\left\{\frac{\left(\sum_{j=1}^{m}\gamma_j C_j\right)}{\left(\sum_{j=1}^{m}\gamma_j\right)^d}\right\}\right\} \tag{2.27}$$

Proof. Suppose that all the $N = kN^*$ servers in the system are indexed by the set $\mathcal{S}_N = \{1, 2, \ldots, N\}$. Under this indexing, we denote the capacity of the n^{th} server in the system by $C_{(n)}$, where $n \in \mathcal{S}_N$ and $C_{(n)} \in \mathcal{C}$. For each job, we define a *selection set* to be the subset of d servers sampled at its arrival instant. We denote by p_A the probability with which the subset $A \subseteq \mathcal{S}_N$ is chosen as the selection set for an arrival. Note that the probabilities p_A, $A \subseteq \mathcal{S}_N$, define the randomized job assignment scheme used. In particular, under the SQ(d) scheme, the probability p_A is non-zero only for subsets $A \subseteq \mathcal{S}_N$ which contain exactly d servers and for each such subset $A \subseteq \mathcal{S}_N$, the probability p_A is given by

$$p_A = \frac{1}{\binom{N}{d}} \tag{2.28}$$

Now we apply condition (1.2) of Corollary 1.1 of [33] to find the stability region of the system under consideration.[5] According to this condition, the system under consideration is stable if the arrival rate λ satisfies

[5] We note that the additional conditions (1.11) and (1.12) of [33]. are automatically satisfied since the interarrival times are exponentially distributed.

$$\varrho^{(N)}(\lambda) = \max_{B \subseteq \mathcal{S}_N} \varrho^{(N)}(B, \lambda) < 1, \tag{2.29}$$

where

$$\varrho^{(N)}(B, \lambda) = \left(\mu \sum_{n \in B} C_{(n)} \right)^{-1} N\lambda \sum_{A \subseteq B} p_A, \tag{2.30}$$

Clearly, for the SQ(d) scheme, $\varrho^{(N)}(B, \lambda)$ is non-zero only when the subset B is composed of at least d servers. From such a set B a subset $A \subseteq B$ of size d can be chosen in $\binom{|B|}{d}$ different ways. Hence, combining (2.28), (2.29), and (2.30) we have

$$\varrho^{(N)}(\lambda) = \max_{B \subseteq \mathcal{S}_N, N \geq |B| \geq d} \varrho^{(N)}(B, \lambda), \tag{2.31}$$

and

$$\varrho^{(N)}(B, \lambda) = \frac{N\lambda}{\mu(N)(N-1)\dots(N-d+1)} \frac{(|B|)(|B|-1)\dots(|B|-d+1)}{\sum_{j \in \mathcal{J}} B_j C_j} < 1, \tag{2.32}$$

where for each $j \in \mathcal{J}$, B_j denotes the number of servers of type j in the set B, and $|B| = \sum_{j \in \mathcal{J}} B_j \geq d$ denotes the cardinality of the set B. Thus, to obtain $\varrho^{(N)}(\lambda)$, we have to maximize $\varrho^{(N)}(B, \lambda)$ over all subsets $B \subseteq \mathcal{S}_N$ having size at least d. This maximization can be done in two steps. The first step is to maximize $\varrho^{(N)}(B, \lambda)$ for given size $|B|$ of the set B by appropriately choosing B_j for all $j \in \mathcal{J}$. The second step is to maximize the the maximum obtained from the first step over all possible values of $|B|$.

Step 1: Clearly, for a fixed size $|B|$ satisfying $N \sum_{j=1}^{m} \gamma_j \leq |B| \leq N \sum_{j=1}^{m+1} \gamma_j$, the term $(|B|)(|B|-1)\dots(|B|-d+1)/\sum_{j \in \mathcal{J}} B_j C_j$, appearing in (2.32), is maximized if $B_j = N\gamma_j$ for $1 \leq j \leq m$, $B_{m+1} = |B| - N \sum_{j=1}^{m} \gamma_j$, and $B_j = 0$ for $j > m+1$.[6] In other words, for any given size $|B|$, the quantity $\varrho^{(N)}(B, \lambda)$ is maximized by filling up the set B starting with lowest capacity servers (with capacity C_1) and gradually moving to servers having higher capacities. The process is continued until the set B is completely filled with $|B|$ servers.

Step 2: Now we shall vary $|B|$ from d to N. We first note that the function $B_1(B_1 - 1)\dots(B_1 - d + 1)/(B_1 C_1)$ is is increasing in B_1. Hence, in the range $d \leq |B| \leq N\gamma_1$,

[6] We have implicitly assumed, without loss of generality, that $N(\min_{j \in \mathcal{J}} \gamma_j) \geq d$, i.e., there are at least d servers of each type. If this is not true, then we can always choose a larger N.

the quantity $\varrho^{(N)}(B,\lambda)$ is maximized at $|B| = B_1 = N\gamma_1$. For $N\gamma_1 \le |B| \le N\gamma_1 + N\gamma_2$, in order to maximize $\varrho^{(N)}(B,\lambda)$ we must have $B_1 = N\gamma_1$ and $B_2 = |B| - N\gamma_1$. Now to have $(N\gamma_1 + B_2)(N\gamma_1 + B_2 - 1)\dots(N\gamma_1 + B_2 - d + 1)/(N\gamma_1 C_1 + B_2 C_2) \le (N\gamma_1 + B_2 + 1)(N\gamma_1 + B_2)\dots(N\gamma_1 + B_2 - d + 2)/(N\gamma_1 C_1 + (B_2+1)C_2)$ we must have $N\gamma_1(C_2 - dC_1) \le (B_2+1)C_2(d-1)$. Hence, as B_2 is increased from 0 to $N\gamma_2$, the expression $(N\gamma_1 + B_2)(N\gamma_1 + B_2 - 1)\dots(N\gamma_1 + B_2 - d + 1)/(N\gamma_1 C_1 + B_2 C_2)$ increases monotonically if $N\gamma_1(C_2 - dC_1) \le C_2(d-1)$; decreases monotonically if $N\gamma_1(C_2 - dC_1) > N\gamma_2 C_2(d-1)$; first decreases and then increases if $N\gamma_1(C_2 - dC_1) > C_2(d-1)$ and $N\gamma_1(C_2 - dC_1) \le N\gamma_2 C_2(d-1)$. Therefore, in the range $d \le |B| \le N(\gamma_1 + \gamma_2)$, the maximum of $\varrho^{(N)}(B,\lambda)$ is obtained either at $|B| = N\gamma_1$ or at $|B| = N(\gamma_1 + \gamma_2)$. Proceeding in this way we find that the maximum of $\varrho^{(N)}(B,\lambda)$ is obtained over $d \le |B| \le N$ at one of the values of $|B|$ in the set $\left\{N\gamma_1, N(\gamma_1+\gamma_2), \dots, N\sum_{j=1}^m \gamma_m, \dots, N\right\}$. Hence,

$$\varrho^{(N)}(\lambda) = \max_{1\le m\le M} \left\{ \frac{\lambda}{\mu \sum_{j=1}^m \gamma_j C_j} \frac{(N\sum_{j=1}^m \gamma_j)(N\sum_{j=1}^m \gamma_j - 1)\dots(N\sum_{j=1}^m \gamma_j - (d-1))}{(N)(N-1)\dots(N-(d-1))} \right\} \tag{2.33}$$

Since $(N\alpha - k)/(N-k)$ is an increasing function in N for $\alpha \le 1$, we have from (2.33) that $\varrho^{(N)}(\lambda)$ increases with N. This implies that $\Lambda_1 \supseteq \Lambda_2 \supseteq \dots$ holds.

We also note that $\frac{N\alpha - k}{N-k} \le \alpha$ for $\alpha \le 1$. Hence, we have from (2.33)

$$\varrho^{(N)}(\lambda) \le \max_{1\le m\le M} \left\{ \frac{\lambda}{\mu \sum_{j=1}^m \gamma_j C_j} (\sum_{j=1}^m \gamma_j)^d \right\} \text{ for all } N. \tag{2.34}$$

The above implies that $\Lambda_\infty \subseteq \Lambda_k$ for all k. Now since $\Lambda_1 \supseteq \Lambda_2 \supseteq \dots$, we have $\Lambda_\infty \subseteq \cap_{k=1}^\infty \Lambda_k$.

Now let us assume $\lambda \in \cap_{k=1}^\infty \Lambda_k$. This implies that $\varrho^{(N)}(\lambda) < 1$ for all N. Hence, $\lim_{N\to\infty} \varrho^{(N)}(\lambda) \le 1$. But from (2.33) we have

$$\lim_{N\to\infty} \varrho^{(N)}(\lambda) = \max_{1\le m\le M} \left\{ \frac{\lambda}{\mu \sum_{j=1}^m \gamma_j C_j} (\sum_{j=1}^m \gamma_j)^d \right\}. \tag{2.35}$$

Hence, we must have $\lambda \in \bar{\Lambda}_\infty$. This proves $\cap_{k=1}^\infty \Lambda_k \subseteq \bar{\Lambda}_\infty$. □

The following result follows from the proof of the above theorem.

Proposition 2.4.1. *If for every $m \in \mathcal{J}$ we have $(\sum_{j=1}^{m-1} \gamma_j)C_m - d(\sum_{j=1}^{m-1} \gamma_j C_j) \leq 0$, then the SQ(d) scheme attains the maximal stability region Λ for heterogeneous systems. In particular, if $d \geq \left\lceil \frac{C_M}{C_1} \right\rceil$, then SQ(d) scheme attains the maximal stability region Λ for heterogeneous systems.*

Proof. Let us assume that for each $m \in \mathcal{J}$ the condition $(\sum_{j=1}^{m-1} \gamma_j)C_m - d(\sum_{j=1}^{m-1} \gamma_j C_j) \leq 0$ holds. From the proof of Theorem 2.4.1 we know that in the range $N \sum_{j=1}^{m-1} \gamma_j \leq |B| \leq N \sum_{j=1}^{m} \gamma_j$ the quantity $\varrho^{(N)}(B, \lambda)$ is maximized by choosing $B_j = N\gamma_j$ for $1 \leq j \leq m-1$, $B_m = |B| - N \sum_{j=1}^{m-1} \gamma_j$, and $B_j = 0$ for $j > m$. Now the term $(|B|)(|B|-1)\dots(|B|-d+1)/(N \sum_{j=1}^{m-1} \gamma_j C_j + B_m C_m)$, appearing in the expression of $\varrho^{(N)}(B, \lambda)$, increases with increase in B_m if

$$\begin{aligned}
&\frac{(N \sum_{j=1}^{m-1} \gamma_j + B_m)(N \sum_{j=1}^{m-1} \gamma_j + B_m - 1)\dots(N \sum_{j=1}^{m-1} \gamma_j + B_m - d + 1)}{N \sum_{j=1}^{m-1} \gamma_j C_j + B_m C_m} \\
&\quad \leq \frac{(N \sum_{j=1}^{m-1} \gamma_j + B_m + 1)(N \sum_{j=1}^{m-1} \gamma_j + B_m)\dots(N \sum_{j=1}^{m-1} \gamma_j + B_m - d + 2)}{N \sum_{j=1}^{m-1} \gamma_j C_j + (B_m + 1) C_m}, \quad (2.36)
\end{aligned}$$

which is equivalent to the condition $(N \sum_{j=1}^{m-1} \gamma_j)C_m - dN(\sum_{j=1}^{m-1} \gamma_j C_j) \leq (B_m + 1)(d - 1)C_m$. But the above condition is true for all $B_{m+1} \geq 0$ since by assumption we have $(\sum_{j=1}^{m-1} \gamma_j)C_m - d(\sum_{j=1}^{m-1} \gamma_j C_j) \leq 0$. Hence, in the range $N \sum_{j=1}^{m-1} \gamma_j \leq |B| \leq N \sum_{j=1}^{m} \gamma_j$ the quantity $\varrho^{(N)}(B, \lambda)$ is monotonically non-decreasing. Since, this true for all $m \in \mathcal{J}$, we conclude that the maximum of $\varrho^{(N)}(B, \lambda)$ in the entire range $d \leq |B| \leq N$ is attained at $|B| = N$. Hence, from (2.33) we have

$$\begin{aligned}
\varrho^{(N)}(\lambda) &= \max_{1 \leq m \leq M} \left\{ \frac{\lambda}{\mu \sum_{j=1}^{m} \gamma_j C_j} \frac{(N \sum_{j=1}^{m} \gamma_j)(N \sum_{j=1}^{m} \gamma_j - 1)\dots(N \sum_{j=1}^{m} \gamma_j - (d-1))}{(N)(N-1)\dots(N-(d-1))} \right\} \\
&= \frac{\lambda}{\mu \sum_{j=1}^{m} \gamma_j C_j}.
\end{aligned}$$

The above implies $\Lambda_k = \Lambda$ for all k and thus proves the first statement of the proposition. The second statement follows easily from the first since, for $d \geq C_M/C_1$, the condition $(\sum_{j=1}^{m-1} \gamma_j)C_m - d(\sum_{j=1}^{m-1} \gamma_j C_j) \leq 0$ is satisfied for all $m \in \mathcal{J}$. □

Remark 6. The above proposition shows that by increasing the number of choices d, the stability region under the SQ(d) scheme can be made equal to the maximal stability region. In particular, if d is chosen such that $d \geq \lceil C_M/C_1 \rceil$, then the maximal stability region can be recovered.

2.5 Analysis of the hybrid SQ(d) scheme

We saw that under the SQ(d) scheme, the stability region of the heterogeneous system is (in general) smaller than the maximum achievable stability region. In this section, we show that with the hybrid SQ(d) scheme it is possible to recover the maximal stability region Λ.

In the hybrid SQ(d) scheme, a server type $j \in \mathcal{J}$ is chosen for a new arrival with a probability p_j independent of all other server types. Hence, for each $j \in \mathcal{J}$, the aggregate arrival process to the set of $N\gamma_j$ servers of type j is Poisson with rate $p_j N\lambda$ and is independent of the arrival processes to other server types. Thus, under this scheme, the system can be viewed as a collection of M independent homogeneous subsystems, working in parallel. The j^{th} subsystem has $N\gamma_j$ identical servers of capacity C_j; the arrival rate at this subsystem is $N\lambda p_j$; and the arriving jobs are assigned to servers in this subsystem according to the SQ(d) scheme. From the results discussed in Section 2.3.2, we know that the j^{th} subsystem is stable if $\rho_j = p_j N\lambda/\mu N\gamma_j C = p_j\lambda/\gamma_j\mu C_j < 1$. Hence, the entire system is stable if $\rho_j < 1$ for all $j \in \mathcal{J}$.

Define $\vec{p} = (p_j, j \in \mathcal{J})$ to be the vector of routing probabilities or the *routing vector*. We call a routing vector $\vec{p}$ stable for a given λ, if the system is stable under the hybrid SQ(d) scheme with $\vec{p}$ as the routing vector and λ as the arrival rate. We first find the necessary and sufficient condition (on λ) that guarantees the existence of a stable routing vector.

Proposition 2.5.1. *There exists a stable routing vector $\vec{p} = (p_j, j \in \mathcal{J})$ for a given λ if and only if $\lambda \in \Lambda$.*

Proof. Let us assume that $\lambda \in \Lambda$ holds. We choose the routing probabilities as

$$p_i = \frac{\gamma_i C_i}{\sum_{j\in\mathcal{J}} \gamma_j C_j}, \text{ for } i \in \mathcal{J} \tag{2.37}$$

For the above choice of p_i, we have

$$\begin{aligned}\rho_i &= \frac{p_i\lambda}{\gamma_i\mu C_i}\\ &= \frac{\gamma_i C_i}{\sum_{j\in\mathcal{J}}\gamma_j C_j}\frac{\lambda}{\gamma_i\mu C_i}\\ &= \frac{\lambda}{\mu\sum_{j\in\mathcal{J}}\gamma_j C_j} < 1 \text{ since } \lambda\in\Lambda\end{aligned}$$

Since the above holds for all $i\in\mathcal{J}$, the system is stable for the above choice of routing probabilities. Hence, the condition $\lambda\in\Lambda$ is sufficient for the existence of routing probabilities for which the system is stable.

We now show that $\lambda\in\Lambda$ is also a necessary condition for existence of a stable routing vector. Let us assume that the system is stable, i.e., $\rho_j<1$ for all $j\in\mathcal{J}$ but $\lambda\notin\Lambda$, i.e., $\lambda/(\mu\sum_{j=1}^{M}\gamma_j C_j)\geq 1$. We show that this leads to contradiction. From the condition above, we have $\lambda/(\mu\sum_{j=1}^{M}\rho_j\gamma_j C_j)>1$ since $\rho_j<1$. We know that the routing probability p_i is related to ρ_i as follows:

$$p_i = \frac{\rho_i\gamma_i\mu C_i}{\lambda} \tag{2.38}$$

Hence, we have

$$\sum_{i\in\mathcal{J}} p_i = \mu\sum_{i\in\mathcal{J}}\frac{\rho_i\gamma_i C_i}{\lambda} < 1, \tag{2.39}$$

which contradicts the fact that $\sum_{i\in\mathcal{J}} p_i = 1$ Hence, the condition $\lambda\in\Lambda$ is necessary for existence of a stable routing vector. □

Remark 7. We have thus shown that for all arrival rates $\lambda\in\Lambda$ it is possible to find a stable routing vector. In particular, from the above proof it is clear that if the routing probabilities are chosen as

$$p_i = \frac{\gamma_i C_i}{\sum_{j\in\mathcal{J}}\gamma_j C_j} \text{ for all } i\in\mathcal{J}, \tag{2.40}$$

then the system is stable under the hybrid SQ(d) scheme for all $\lambda\in\Lambda$. We note that the above choice of the routing probabilities requires the knowledge of the system parameters (e.g., the knowledge of $\gamma_j, C_j,\ j\in\mathcal{J}$). In general, the routing probabilities for which the system operating under the hybrid SQ(d) scheme is stable are functions of the system parameters.

Henceforth we shall assume that $\lambda \in \Lambda$ holds. Therefore, the existence of a stable routing vector $\vec{p}$ is guaranteed. We now compute the routing vector $\vec{p}^* = (p_j^*, j \in \mathcal{J})$ or equivalently the load vector $\vec{\rho}^* = (\rho_j^*, j \in \mathcal{J})$ that minimizes the mean sojourn time of jobs in the limiting system under the hybrid SQ(d) scheme for a given value of $\lambda \in \Lambda$.

From the discussion in Section 2.3.2, we know that for a given routing vector $\vec{p} = (p_j, j \in \mathcal{J})$, the mean sojourn time of jobs in the j^{th} subsystem (in the limit as $N \to \infty$) is given by

$$\bar{T}_j^{\text{hsqd}}(\vec{\rho}) = \frac{\gamma_j}{p_j \lambda} \sum_{n=1}^{\infty} \rho_j^{\frac{d^n-1}{d-1}}, \tag{2.41}$$

where $\vec{\rho} = (\rho_j, j \in \mathcal{J})$ and $\rho_j = p_j \lambda / \mu \gamma_j C_j$. Hence, the overall mean sojourn time of jobs in the system is given by

$$\bar{T}^{\text{hsqd}}(\vec{\rho}) = \sum_{j \in \mathcal{J}} p_j \bar{T}_j^{hsqd}(\vec{\rho}) = \frac{1}{\lambda} \sum_{j \in \mathcal{J}} \gamma_j \sum_{n=1}^{\infty} \rho_j^{\frac{d^n-1}{d-1}}. \tag{2.42}$$

We now formulate the mean sojourn time minimization problem in terms of the loads ρ_j, $j \in \mathcal{J}$, as follows:

$$\begin{aligned} \underset{\vec{\rho}}{\text{Minimize}} \quad & \frac{1}{\lambda} \sum_{j \in \mathcal{J}} \gamma_j \sum_{n=1}^{\infty} \rho_j^{\frac{d^n-1}{d-1}} \\ \text{subject to} \quad & 0 \le \rho_j \le 1, \text{for all } j \in \mathcal{J} \\ & \sum_{j \in \mathcal{J}} \gamma_j C_j \rho_j = \frac{\lambda}{\mu}. \end{aligned} \tag{2.43}$$

Note that the equality constraint in (2.43) ensures that the routing probabilities sum to unity. To characterize the solution of the convex problem defined in (2.43), we proceed along lines similar to Theorem 1 of [36]. Let $\mathcal{J}_{\text{opt}} \subseteq \mathcal{J}$ denote the index set of server types being used in the optimal scheme.

Proposition 2.5.2. *Let $\Phi : \mathbb{R}_+ \to [0,1)$ be the inverse of the monotone mapping $\Phi^{-1} : [0,1) \to \mathbb{R}_+$ defined as $\Phi^{-1}(\rho) = \sum_{k=1}^{\infty} (d^k - 1) \rho^{\frac{d^k-d}{d-1}} / (d-1) < \sum_{x=1}^{\infty} x \rho^{x-1} < \infty$ for*

$0 < \rho < 1$. *Further, for each* $j \in \mathcal{J}$, *let* $\Psi_j : \mathbb{R}_+ \to \mathbb{R}_+$ *denote the inverse of the monotone mapping* $\Psi_j^{-1} : \mathbb{R}_+ \to \mathbb{R}_+$ *defined as* $\Psi_j^{-1}(\theta) = \mu \sum_{i=j}^{M} \gamma_i C_i \Phi(\theta C_i)$. *The index set of server types used in the hybrid SQ(d) scheme is then given by* $\mathcal{J}_{opt} = \{j^*, j^*+1, \ldots, M\}$, *where* j^* *is given by*

$$j^* = \inf\left\{ j \in \mathcal{J} : \frac{1}{C_j} < \Psi_j(\lambda) \right\}. \tag{2.44}$$

Moreover, the optimal traffic intensities ρ_i^*, *for* $i \in \mathcal{J}$ *satisfy*

$$\rho_i^* = \begin{cases} \Phi(\Psi_{j*}(\lambda) C_i), & \text{if } i \in \mathcal{J}_{opt} \\ 0, & \text{otherwise.} \end{cases} \tag{2.45}$$

Proof. The Lagrangian associated with problem (2.43) is given by

$$L(\vec{\rho}, \vec{\nu}, \vec{\zeta}, \theta) = \sum_{j=1}^{M} \gamma_j \sum_{k=1}^{\infty} \rho_i^{\frac{d^k-1}{d-1}} + \sum_{j=1}^{M} \nu_j (0 - \rho_j) + \sum_{j=1}^{M} \zeta_j (\rho_j - 1) + \theta \left(\sum_{j=1}^{M} \gamma_j C_j \rho_j - \frac{\lambda}{\mu} \right), \tag{2.46}$$

where $\vec{\nu} = (\nu_1, \nu_2, \ldots, \nu_M) \geq \vec{0}$, $\vec{\zeta} = (\zeta_1, \zeta_2, \ldots, \zeta_M) \geq \vec{0}$, and $\theta \in \mathbb{R}$. Since problem (2.43) is strictly convex and a feasible solution exits (due to condition $\lambda \in \Lambda$), by Slater's condition [[37], Section 5.2.3], strong duality is satisfied. Let $\vec{\rho}^*$ and $\left(\vec{\nu}^*, \vec{\zeta}^*, \theta^*\right)$ denote the primal and dual optimal solutions, respectively. Since the duality gap is zero, we have by applying KKT conditions

$$\vec{0} \leq \vec{\rho}^* \leq \vec{1}$$

$$\sum_{j=1}^{M} \gamma_j C_j \rho_j^* = \frac{\lambda}{\mu}$$

$$\theta^* \in \mathbb{R}, \vec{\nu}^* \geq \vec{0}, \vec{\zeta}^* \geq \vec{0}$$

$$\nu_j^* \rho_j^* = 0, \zeta_j^* (\rho_j^* - 1) = 0 \ \forall j \in \mathcal{J} \tag{2.47}$$

$$\gamma_j \sum_{k=1}^{\infty} \frac{d^k - 1}{d-1} \left(\rho_j^*\right)^{\frac{d^k - d}{d-1}} - \theta^* \gamma_j C_j - \nu_j^* + \zeta_j^* = 0 \ \forall j \in \mathcal{J}. \tag{2.48}$$

Since the objective function tends to infinity as $\rho_j^* \to 1$, for any $j \in \mathcal{J}$, we must have $\vec{\rho}^* < \vec{1}$ which implies $\vec{\zeta}^* = \vec{0}$ from (2.47). Since $\vec{\nu}^* \geq \vec{0}$, we have from (2.48)

$$\theta^* \leq \frac{1}{C_j} \sum_{k=1}^{\infty} \frac{d^k - 1}{d-1} \left(\rho_j^*\right)^{\frac{d^k - d}{d-1}} \quad \forall j \in \mathcal{J} \tag{2.49}$$

Further, by eliminating ν_j^* from (2.48) we obtain

$$\left(\sum_{k=1}^{\infty} \frac{d^k - 1}{d-1} \left(\rho_j^*\right)^{\frac{d^k - d}{d-1}} - \theta^* C_j \right) \rho_j^* = 0 \tag{2.50}$$

Thus, if for a given $j \in \mathcal{J}$, we have $\theta^* > \frac{1}{C_j}$, then (2.49) implies for that j we must have $\rho_j^* > 0$. Therefore, from (2.50) and from the definition of the map Φ we have $\rho_j^* = \Phi(\theta^* C_j)$ for such a j. If, on the other hand, $\theta^* \leq \frac{1}{C_j}$ for some $j \in \mathcal{J}$, then $\rho_j^* = 0$. This is so because $\theta^* \leq \frac{1}{C_j}$ implies $\theta^* < \frac{1}{C_j} \sum_{k=1}^{\infty} \frac{d^k-1}{d-1} \left(\rho_j^*\right)^{\frac{d^k-d}{d-1}}$ if $\rho_j^* > 0$. But this contradicts (2.50). Hence, we must have

$$\rho_j^* = \begin{cases} \Phi(\theta^* C_j), & \text{if } \frac{1}{C_j} < \theta^* \\ 0, & \text{otherwise.} \end{cases} \tag{2.51}$$

To find θ^*, we use the equality constraint in (2.43). If the server types belonging to the set $\{j^*, j^*+1, \ldots, M\}$ are used in the optimal SQ(d) scheme, then

$$\sum_{j=j^*}^{M} \gamma_j C_j \Phi(\theta^* C_j) = \frac{\lambda}{\mu} \tag{2.52}$$

Hence by definition of the map Ψ_j,

$$\theta^* = \Psi_{j*}(\lambda), \tag{2.53}$$

where j^* is defined as in (2.44). □

Remark 8. We note that the optimal routing probabilities are functions of the arrival rate λ. Hence, computing them requires the knowledge of the arrival rate of jobs, which is difficult to estimate. However, as discussed in Remark 7 the knowledge of only the system parameters is required to compute the routing probabilities which stabilize the system for all $\lambda \in \Lambda$. Such a choice of the routing probabilities is, however, not optimal in terms of minimizing the mean sojourn time of jobs in the system.

2.6 Analysis of the type-based scheme

In the previous section, we saw that the hybrid SQ(d) scheme achieves the maximal stability region only when the proportions γ_j, $j \in \mathcal{J}$ are known. Hence, if a large number servers of a certain type fail or need to be shut down (due to low utilization), then a reconfiguration of the job dispatcher is required. In this section, we show that the type-based scheme achieves the stability region without requiring the knowledge of the proportions γ_j, $j \in \mathcal{J}$. Therefore, the type-based scheme is more robust to server failures than the hybrid SQ(d) scheme. We also characterize the performance of the type-based scheme in the large system limit using mean field analysis.

As in the SQ(d) scheme, we denote by $x_{k,j}^{(N)}(t)$, $k \in \mathbb{Z}_+, j \in \mathcal{J}, t \geq 0$ the fraction of type j servers having at least k unfinished jobs at time t in the finite system of size N and define the process $\mathbf{x}^{(N)}(\cdot) = (x_{k,j}^{(N)}(\cdot), k \in \mathbb{Z}_+, j \in \mathcal{J})$. Clearly, the process $\mathbf{x}^{(N)}(\cdot)$ takes values in the space $\mathcal{U}^{(N)}$. Moreover, under the assumptions of Poisson arrivals and exponential job length distribution, the process $\mathbf{x}^{(N)}(\cdot)$ is Markov. We first find the set of arrival rates for which the process $\mathbf{x}^{(N)}(\cdot)$ is positive recurrent. This set is the stability region of the system working under the type-based scheme.

Theorem 2.6.1. *The system under consideration is stable under the type-based scheme if $\lambda \in \Lambda$.*

Proof. We proceed along lines similar to the proof of Theorem 2.4.1. We first index the N servers in the system by the set $\mathcal{S}_N = \{1, 2, \ldots, N\}$. The capacity of the n^{th} server is denoted by $C_{(n)}$, where $C_{(n)} \in \mathcal{C}$ and $n \in \mathcal{S}_N$. For each job, we define a *selection set* to be the subset of $\sum_{j \in \mathcal{J}} d_j$ servers sampled at its arrival. We denote by p_A the probability that the subset $A \subseteq \mathcal{S}_N$ is chosen as the selection set for an arrival. Under the type-based scheme, the probability p_A is non-zero only for subsets A which contain d_j servers of type j for all $j \in \mathcal{J}$ and for each such a subset A, the probability p_A is given by

$$p_A = \frac{1}{\prod_{j \in \mathcal{J}} \binom{N\gamma_j}{d_j}}. \tag{2.54}$$

Applying condition (1.2) of Corollary 1.1 of [33], we have that the system under consideration is stable if the arrival rate λ satisfies

$$\varrho^{(N)}(\lambda) = \max_{B \subseteq \mathcal{S}_N} \varrho^{(N)}(B, \lambda) < 1, \tag{2.55}$$

where

$$\varrho^{(N)}(B,\lambda) = \left(\mu \sum_{n\in B} C_{(n)}\right)^{-1} N\lambda \sum_{A\subseteq B} p_A. \tag{2.56}$$

Clearly, for the type based scheme, the $\varrho^{(N)}(B,\lambda)$ is non-zero only when the subset B is composed of at least d_j servers of type j for all $j \in \mathcal{J}$. Let B_j ($\geq d_j$) denote the number of type j servers in B. Hence, a set A having exactly d_j servers of type j for all $j \in \mathcal{J}$ can be chosen from the set B in $\prod_{j\in\mathcal{J}} \binom{B_j}{d_j}$ different ways. Hence, using (2.54), (2.55), and (2.56) we have

$$\varrho^{(N)}(\lambda) = \max_{B\subseteq \mathcal{S}_N : N\gamma_j \geq B_j \geq d_j \forall\, j\in\mathcal{J}} \varrho^{(N)}(B,\lambda), \tag{2.57}$$

and

$$\varrho^{(N)}(B,\lambda) = \frac{N\lambda}{\mu} \frac{1}{\sum_{j\in\mathcal{J}} B_j C_j} \prod_{j\in\mathcal{J}} \frac{\binom{B_j}{d_j}}{\binom{N\gamma_j}{d_j}} \tag{2.58}$$

It is easy to verify that that the function $\frac{\prod_{j\in\mathcal{J}}\binom{B_j}{d_j}}{\sum_{j\in\mathcal{J}} B_j C_j}$ is increasing with respect to B_j for each $j \in \mathcal{J}$. Hence, $\varrho^{(N)}(B,\lambda)$ is maximized when we set $B_j = N\gamma_j$ for all $j \in \mathcal{J}$. Thus, we have

$$\varrho^{(N)}(\lambda) = \frac{N\lambda}{\mu} \frac{1}{N\sum_{j\in\mathcal{J}} \gamma_j C_j} = \frac{\lambda}{\mu \sum_{j\in\mathcal{J}} \gamma_j C_j} \tag{2.59}$$

This implies that the system under consideration is stable if $\lambda \in \Lambda$. □

Remark 9. Thus, Theorem 2.6.1 shows that the system under consideration is stable for all N under the type-based scheme if $\lambda \in \Lambda$. We note that above proof does not depend the choices of d_j, $j \in \mathcal{J}$. Hence, even with $d_j = 1$ for all $j \in \mathcal{J}$ the type-based scheme achieves the maximal stability region. Furthermore, the scheme does not require knowledge of the proportions γ_j, $j \in \mathcal{J}$, to achieve the maximal stability region.

Remark 10. An alternative proof of stability via a coupling argument is as follows: Consider a modified scheme in which, upon arrival of each job, one server is chosen from each type uniformly at random (i.e., $d_j = 1$ for all $j \in \mathcal{J}$). The job is then routed to the sampled server of type j with probability $\frac{\gamma_j C_j}{\sum_{i\in\mathcal{J}} \gamma_i C_i}$ for each $j \in \mathcal{J}$. A coupling argument, similar to the one discussed in the proof of Theorem 3 of [30], shows that the system operating under the modified scheme always has higher number of unfinished jobs than that operating under the type-based scheme. It is easy to check that the system operating under the modified scheme is stable for $\lambda \in \Lambda$. Hence, the system operating under the type-based scheme also must be stable under for $\lambda \in \Lambda$.

We now analyze the evolution of the process $\mathbf{x}^{(N)}(\cdot)$ in the limit as $N \to \infty$ using mean field techniques. The generator $\mathbf{A}_N$ of the process $\mathbf{x}^{(N)}(\cdot)$ acting on continuous functions $\varphi : \bar{\mathcal{U}}^M \to \mathbb{R}$ is given by $\mathbf{A}_N \varphi(\mathbf{u}) = \sum_{\mathbf{v} \neq \mathbf{u}} r(\mathbf{u} \to \mathbf{v}) (\varphi(\mathbf{v}) - \varphi(\mathbf{u}))$, where $r(\mathbf{u} \to \mathbf{v})$ 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 $\mathbf{A}_N$.

Lemma 2.6.1. *Let $\mathbf{u} \in \mathcal{U}^{(N)}$ be any state of the process $\mathbf{x}^{(N)}(\cdot)$ and let $\mathbf{e}(n,j) = (e_{k,i}, k \in \mathbb{Z}_+, i \in \mathcal{J})$ be the unit vector with $e_{n,j} = 1$ and $e_{k,i} = 0$ for $(k,i) \neq (n,j)$. Under the type-based scheme, the generator $\mathbf{A}_N$ of the Markov process $\mathbf{x}^{(N)}(\cdot)$ acting on continuous functions $\varphi : \bar{\mathcal{U}}^M \to \mathbb{R}$ is given by*

$$\begin{aligned}\mathbf{A}_N \varphi(\mathbf{u}) = N\lambda \sum_{j \in \mathcal{J}} \sum_{n \geq 1} \left[(u_{n-1,j})^{d_j} - (u_{n,j})^{d_j} \right] \prod_{i=1}^{j-1} (u_{n-1,i})^{d_i} \prod_{i=j+1}^{M} (u_{n,i})^{d_i} \\ \times \left[\varphi\left(\mathbf{u} + \frac{\mathbf{e}(n,j)}{N\gamma_j} \right) - \varphi(\mathbf{u}) \right] \\ + \mu N \sum_{j \in \mathcal{J}} \sum_{n \geq 1} \gamma_j C_j \left[u_{n,j} - u_{n+1,j} \right] \left[\varphi\left(\mathbf{u} - \frac{\mathbf{e}(n,j)}{N\gamma_j} \right) - \varphi(\mathbf{u}) \right]. \quad (2.60)\end{aligned}$$

Proof. We first consider an arrival joining a server of type j having exactly $n-1$ unfinished jobs, when the state of the system is $\mathbf{u}$. This corresponds to a transition from the state $\mathbf{u}$ to the state $\mathbf{u} + \frac{\mathbf{e}(n,j)}{N\gamma_j}$. We note that a job joins a server of type j with exactly $n-1$ occupancy if the following conditions are satisfied:

- Among the d_j sampled servers of type j, at least one has exactly $n-1$ jobs and the rest of them have at least n jobs. This occurs with probability $\left((u_{n-1,j})^{d_j} - (u_{n,j})^{d_j} \right)$.
- For each $i < j$, all the d_i sampled servers of type i have at least $n-1$ jobs. This occurs with probability $\prod_{i=1}^{j-1} (u_{n-1,i})^{d_i}$.
- For each $i > j$, all the d_i servers of type i have at least n jobs. This occurs with probability $\prod_{i=j+1}^{M} (u_{n,i})^{d_i}$.

Since the arrival rate of jobs is $N\lambda$, the rate of the above transition is given by

$$r\left(\mathbf{u} \to \mathbf{u} + \frac{\mathbf{e}(n,j)}{N\gamma_j}\right) = N\lambda\left[(u_{n-1,j})^{d_j} - (u_{n,j})^{d_j}\right]\prod_{i=1}^{j-1}(u_{n-1,i})^{d_i}\prod_{i=j+1}^{M}(u_{n,i})^{d_i} \tag{2.61}$$

Further, the rate at which jobs depart from a server of type j having exactly n jobs is $\mu C_j N\gamma_j\left(u_{n,j} - u_{n+1,j}\right)$. Hence,

$$r\left(\mathbf{u} \to \mathbf{u} - \frac{\mathbf{e}(n,j)}{N\gamma_j}\right) = \mu C_j N\gamma_j\left(u_{n,j} - u_{n+1,j}\right). \tag{2.62}$$

The expression (2.60) now follows directly from the definition of $\mathbf{A}_N$. □

The next theorem shows that the process $\mathbf{x}^{(N)}(\cdot)$ weakly converges to a deterministic process $\mathbf{x}(\cdot)$ as $N \to \infty$.

Theorem 2.6.2. *If $\mathbf{x}^N(0) \Rightarrow \mathbf{u}_0 \in \bar{\mathcal{U}}^M$ as $N \to \infty$, then the process $\mathbf{x}^{(N)}(\cdot) \Rightarrow \mathbf{x}(\cdot) = (x_{k,j}(\cdot), k \in \mathbb{Z}_+, j \in \mathcal{J})$, where the process $\mathbf{x}(\cdot)$ is a deterministic process taking values in the space $\bar{\mathcal{U}}^M$ and is given by the unique solution of the following system of differential equations*

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{2.63}$$

$$\dot{\mathbf{x}}(t) = \mathbf{l}(\mathbf{x}(t)), \tag{2.64}$$

where the mapping $\mathbf{l} : \bar{\mathcal{U}}^M \to \left(\mathbb{R}^{\mathbb{Z}_+}\right)^M$ is given by

$$l_{0,j}(\mathbf{x}) = 0, \text{ for } j \in \mathcal{J}, \tag{2.65}$$

$$\begin{aligned} l_{k,j}(\mathbf{x}) = \frac{\lambda}{\gamma_j}\left((x_{k-1,j})^{d_j} - (x_{k,j})^{d_j}\right)\prod_{i=1}^{j-1}(x_{k-1,i})^{d_i}\prod_{i=j+1}^{M}(x_{k,i})^{d_i} \\ - \mu C_j\left(x_{k,j} - x_{k+1,j}\right), \text{ for } k \geq 1, j \in \mathcal{J}. \end{aligned} \tag{2.66}$$

The process $\mathbf{x}(\cdot)$, defined in the theorem above, is referred to as the *mean field limit* of the heterogeneous system under the type-based scheme. Before proving the above theorem we first show that the process $\mathbf{x}(\cdot)$ defined by (2.63)-(2.64) is indeed unique in the space $\bar{\mathcal{U}}^M$.

Proposition 2.6.1. *If $\mathbf{u}_0 \in \bar{\mathcal{U}}^M$, then the system defined by (2.63)-(2.64) has a unique solution $\mathbf{x}(\cdot)$ taking values in the space $\bar{\mathcal{U}}^M$.*

Proof. Define $\theta(x) = [\min(x, 1)]_+$, where $[z]_+ = \max\{0, z\}$ and let us consider the following modification of (2.63)-(2.64):

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{2.67}$$

$$\dot{\mathbf{x}}(t) = \hat{\mathbf{l}}(\mathbf{x}(t)), \tag{2.68}$$

where the mapping $\hat{\mathbf{l}} : \left(\mathbb{R}^{\mathbb{Z}_+}\right)^M \to \left(\mathbb{R}^{\mathbb{Z}_+}\right)^M$ is given by

$$\hat{l}_{0,j}(\mathbf{x}) = 0, \text{ for } j \in \mathcal{J}, \tag{2.69}$$

$$\hat{l}_{k,j}(\mathbf{x}) = \frac{\lambda}{\gamma_j}\left[\left(\theta\left(x_{k-1,j}\right)\right)^{d_j} - \left(\theta\left(x_{k,j}\right)\right)^{d_j}\right]_+ \prod_{i=1}^{j-1}\left(\theta\left(x_{k-1,i}\right)\right)^{d_i} \tag{2.70}$$

$$\times \prod_{i=j+1}^{M}\left(\theta\left(x_{k,i}\right)\right)^{d_i} - \mu C_j\left[\theta\left(x_{k,j}\right) - \theta\left(x_{k+1,j}\right)\right]_+, \text{ for } k \geq 1, j \in \mathcal{J}.$$

Clearly, the right hand sides of (2.66) and (2.70) are equal if $\mathbf{x} \in \bar{\mathcal{U}}^M$. Therefore, the two systems must have identical solutions in $\bar{\mathcal{U}}^M$. Also if $\mathbf{u}_0 \in \bar{\mathcal{U}}^M$, then any solution of the modified system remains within $\bar{\mathcal{U}}^M$. This is because of the facts that if $x_{n,j}(t) = x_{n+1,j}(t)$ for some j, n, t, then $\hat{l}_{n,j}(\mathbf{x}(t)) \geq 0$ and $\hat{l}_{n+1,j}(\mathbf{u}(t)) \leq 0$; and if $x_{n,j}(t) = 0$ for some j, n, t, then $\hat{l}_{n,j}(\mathbf{x}(t)) \geq 0$. Hence, to prove the uniqueness of solution of (2.63)-(2.64), we need to show that the modified system (2.67)-(2.68) has a unique solution in $(\mathbb{R}^{\mathbb{Z}_+})^M$. We now extend the metric ω defined in (2.5) to the space $(\mathbb{R}^{\mathbb{Z}_+})^M$.

Using (2.5) and the facts that $|x_+ - y_+| \leq |x - y|$ for any $x, y \in \mathbb{R}$, $|a_1 b_1^m - a_2 b_2^m| \leq |a_1 - a_2| + m\,|b_1 - b_2|$ for any $a_1, a_2, b_1, b_2 \in [0, 1]$, and $|\theta(x) - \theta(y)| \leq |x - y|$ for any $x, y \in \mathbb{R}$ we obtain

$$\omega(\hat{\mathbf{l}}(\mathbf{x}), 0) \leq K_1, \tag{2.71}$$

$$\omega(\hat{\mathbf{l}}(\mathbf{x}), \hat{\mathbf{l}}(\mathbf{y})) \leq K_2\,\omega(\mathbf{x}, \mathbf{y}), \tag{2.72}$$

where $\mathbf{x}, \mathbf{y} \in (\mathbb{R}^{\mathbb{Z}_+})^M$, K_1 and K_2 are constants defined as $K_1 = \frac{\lambda}{\min_{j\in\mathcal{J}} \gamma_j} + \mu(\max_{j\in\mathcal{J}} C_j)$ and $K_2 = 4M\lambda\frac{\max_{j\in\mathcal{J}} d_j}{\min_{j\in\mathcal{J}} \gamma_j} + 3\mu(\max_{1\leq j\leq M} C_j)$. The existence and uniqueness of solution of (2.67)-(2.68) now follows from inequalities (2.71) and (2.72) by using Picard's iteration technique since $(\mathbb{R}^{\mathbb{Z}_+})^M$ is complete under ω. $\square$

To prove Theorem 2.6.2, we will also require the following result, which shows that the mean field process $\mathbf{x}(\cdot)$ is smooth with respect to the initial conditions. We denote by $\mathbf{x}(t,\mathbf{u})$ the value taken by the process $\mathbf{x}(\cdot)$ at time $t \geq 0$ when it starts at $\mathbf{x}(0) = \mathbf{u}$.

Lemma 2.6.2. *For each $j, j', i \in \mathcal{J}$, $n, n', k \in \mathbb{Z}_+$, and $t \geq 0$, the partial derivatives $\frac{\partial \mathbf{x}(t,\mathbf{u})}{\partial u_{n,j}}$, $\frac{\partial^2 \mathbf{x}(t,\mathbf{u})}{\partial {u_{n,j}}^2}$, and $\frac{\partial^2 \mathbf{x}(t,\mathbf{u})}{\partial u_{n,j} \partial u_{n',j'}}$ exist for $\mathbf{u} \in \bar{\mathcal{U}}^M$ and satisfy*

$$\left| \frac{\partial x_{k,i}(t,\mathbf{u})}{\partial u_{n,j}} \right| \leq \exp(B_1 t) \tag{2.73}$$

and

$$\left| \frac{\partial^2 x_{k,i}(t,\mathbf{u})}{\partial {u_{n,j}}^2} \right|, \left| \frac{\partial^2 x_{k,i}(t,\mathbf{u})}{\partial u_{n,j} \partial u_{n',j'}} \right| \leq \frac{B_2}{B_1} (\exp(2B_1 t) - \exp(B_1 t)), \tag{2.74}$$

where $B_1 = \frac{2\lambda \sum_{j\in\mathcal{J}} d_j}{\min_{j\in\mathcal{J}} \gamma_j} + 2\mu \left(\max_{j\in\mathcal{J}} C_j\right)$, and $B_2 = \frac{2\lambda\left(\sum_{j\in\mathcal{J}} d_j\right)^2}{\min_{j\in\mathcal{J}} \gamma_j}$.

Proof. Fix j, n, $\mathbf{u}$ and define $\mathbf{x}'(t) = \partial \mathbf{x}(t,\mathbf{u})/\partial u_{n,j}$. If this partial derivative exists, then $\mathbf{x}'(t)$ must satisfy ${x'}_{0,i}(t) = 0$, ${x'}_{k,i}(0) = \delta_{i,j}\delta_{k,n}$, where $\delta_{i,j}$ denotes the Kronecker's delta. Further, by differentiating (2.66) with respect to $u_{n,j}$ we obtain (we omit the argument t to simplify notations)

$$\begin{aligned}
\frac{d{x'}_{k,i}}{dt} = \frac{d_i \lambda}{\gamma_j} &\left((x_{k-1,i})^{d_i - 1} {x'}_{k-1,i} - (x_{k,i})^{d_i-1} {x'}_{k,i} \right) \prod_{s=1}^{i-1} (x_{k-1,s})^{d_s} \prod_{s=i+1}^{M} (x_{k,s})^{d_s} \\
&+ \sum_{l=1}^{i-1} \frac{d_l \lambda}{\gamma_j x_{k-1,l}} \left((x_{k-1,i})^{d_i} - (x_{k,i})^{d_i} \right) \prod_{s=1}^{i-1} (x_{k-1,s})^{d_s} \prod_{s=i+1}^{M} (x_{k,s})^{d_s} \\
&+ \sum_{l=i+1}^{M} \frac{d_l \lambda}{\gamma_j x_{k,l}} \left((x_{k-1,i})^{d_i} - (x_{k,i})^{d_i} \right) \prod_{s=1}^{i-1} (x_{k-1,s})^{d_s} \prod_{s=i+1}^{M} (x_{k,s})^{d_s} \\
&- \mu C_j \left({x'}_{k,i} - {x'}_{k+1,i} \right).
\end{aligned} \tag{2.75}$$

Conversely, if $\mathbf{x}'(t)$ is a solution of the system above, then it must be the required partial derivative. Using Lemma 3.1 of [30] (a generalized version of the Gronwall's lemma) with $a = B_1$, $b_0 = 0$, and $c = 1$ and the fact that $|x_{k,i}| \leq 1$ for all k,i it is easily seen that $\frac{\partial x_{k,i}(t,\mathbf{u})}{\partial u_{n,j}}$ exists and is bounded as given by (2.73).

Similarly, by differentiating (2.75) again with respect to $u_{n,j}$ and $u_{n',j'}$, we obtain the systems of equations satisfied by $\frac{\partial^2 x_{k,i}(t,\mathbf{u})}{\partial {u_{n,j}}^2}$ and $\frac{\partial^2 x_{k,i}(t,\mathbf{u})}{\partial u_{n,j} \partial u_{n',j'}}$, respectively. Lemma 3.1 of [30] can be applied again to these systems to show that the second order partial derivatives also exist and are bounded as given by (2.74). □

We now prove Theorem 2.6.2 below.

Proof of Theorem 2.6.2 : Let Ξ be the set of continuous functions $\varphi : \bar{\mathcal{U}}^M \to \mathbb{R}$ and let D be the set of those $\varphi \in \Xi$ for which the derivatives $\frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}}$, $\frac{\partial^2 \varphi(\mathbf{u})}{\partial u_{n,j}^2}$, and $\frac{\partial^2 \varphi(\mathbf{u})}{\partial u_{n,j} \partial u'_{n',j'}}$ exist for all $n, n' \in \mathbb{Z}_+$ and $j, j' \in \mathcal{J}$ and are uniformly bounded by some constant $B < \infty$. Using the metric ω on $\bar{\mathcal{U}}^M$ and the sup norm on Ξ we find that D is dense in Ξ. For $\varphi \in D$ we have

$$N\gamma_j \left(\varphi \left(\mathbf{u} + \frac{\mathbf{e}(n,j)}{N\gamma_j} \right) - \varphi(\mathbf{u}) \right) \to \frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}} \tag{2.76}$$

$$N\gamma_j \left(\varphi \left(\mathbf{u} - \frac{\mathbf{e}(n,j)}{N\gamma_j} \right) - \varphi(\mathbf{u}) \right) \to -\frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}}. \tag{2.77}$$

Thus using (2.60) we have

$$\mathbf{A}_N \varphi(\mathbf{u}) \to \sum_{j \in \mathcal{J}} \sum_{n \geq 1} \frac{\lambda}{\gamma_j} \left[(u_{n-1,j})^{d_j} - (u_{n,j})^{d_j} \right] \prod_{i=1}^{j-1} (u_{n-1,i})^{d_i} \prod_{i=j+1}^{M} (u_{n,i})^{d_i} \left(\frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}} \right)$$
$$- \mu \sum_{j \in \mathcal{J}} \sum_{n \geq 1} C_j \left[u_{n,j} - u_{n+1,j} \right] \left(\frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}} \right). \tag{2.78}$$

The right hand side of (2.78) can be rewritten as

$$\sum_{j \in \mathcal{J}} \sum_{n \geq 1} \left(\frac{\lambda}{\gamma_j} \left[(u_{n-1,j})^{d_j} - (u_{n,j})^{d_j} \right] \prod_{i=1}^{j-1} (u_{n-1,i})^{d_i} \prod_{i=j+1}^{M} (u_{n,i})^{d_i} - \mu C_j \left(u_{n,j} - u_{n+1,j} \right) \right)$$
$$\times \left(\frac{\partial \varphi(\mathbf{u})}{\partial u_{n,j}} \right), \tag{2.79}$$

which coincides with

$$\frac{d}{dt} \varphi(\mathbf{x}(t, \mathbf{u}))|_{t=0}, \tag{2.80}$$

where $\mathbf{x}(t,\mathbf{u})$ is the solution of (2.63)-(2.64) with $\mathbf{x}(0) = \mathbf{u}$.

We know that the semigroups of operators $(\mathbf{T}(t), t \geq 0)$ and $(\mathbf{T}_N(t), t \geq 0)$ corresponding to the processes $\mathbf{x}(\cdot)$ and $\mathbf{x}^{(N)}(\cdot)$ are given by

$$\mathbf{T}(t)\varphi(\mathbf{u}) = \varphi(\mathbf{x}(t,\mathbf{u})), \tag{2.81}$$

$$\mathbf{T}_N(t)\varphi(\mathbf{u}) = \mathbb{E}[\varphi(\mathbf{x}^{(N)}(t))|\mathbf{x}^{(N)}(0) = \mathbf{u}]. \tag{2.82}$$

The generators corresponding to the semigroups $\mathbf{T}$ and $\mathbf{T}_N$ are $\mathbf{A}$ and $\mathbf{A}_N$, respectively, where

$$\mathbf{A}\varphi(\mathbf{u}) = \frac{d}{dt}\varphi(\mathbf{x}(t,\mathbf{u}))|_{t=0}, \tag{2.83}$$

and $\mathbf{A}_N$ is given by (2.60). Hence, from (2.78),(2.79)(2.80) we have

$$\lim_{N\to\infty} \mathbf{A}^{(N)}\varphi = \mathbf{A}\varphi \tag{2.84}$$

for all $\varphi \in D$.

Define $D_0 \subset D$ as the set of those functions in D which depend only on finitely many components $u_{n,j}$. By definition of the metric ω on $\bar{\mathcal{U}}^M$, D_0 is dense in D and hence in Ξ. Also, it follows from Lemma 2.6.2 that $\mathbf{T}(t)\varphi_0 \in D$ for $\varphi_0 \in D_0$ and $t \geq 0$. Therefore, by Proposition C.2.1 we have that D is the core of A. We also observe that the semigroups $(\mathbf{T}_N(t), t \geq 0)$ and $(\mathbf{T}(t), t \geq 0)$ are, by definition, strongly continuous, contraction semigroups on Ξ. These facts together with (2.84) and Theorem C.2.3 imply that $\mathbf{T}_N(t)\varphi \to \mathbf{T}(t)\varphi$ for all $\varphi \in \Xi$ and all $t \geq 0$.

Now we notice that $\mathbf{T}$ is a Feller semigroup on Ξ. This is because i) $\mathbf{T}(t)1 = 1$, where 1 is the indicator function on $\bar{\mathcal{U}}^M$, ii) by Lemma 2.6.2, $\mathbf{x}(t,\mathbf{u})$ is continuous with respect to initial condition $\mathbf{u}$. Hence, applying Theorem C.2.2 we conclude that if $\mathbf{x}^N(0) \Rightarrow \mathbf{u}_0 \in \bar{\mathcal{U}}^M$, then $\mathbf{x}^{(N)}(\cdot) \Rightarrow \mathbf{x}(\cdot)$. □

Now we characterize the properties of the equilibrium points of the mean field $\mathbf{x}(\cdot)$. We recall that a point $\mathbf{P} = (P_{k,j}, k \in \mathbb{Z}_+, j \in \mathcal{J}) \in \bar{\mathcal{U}}^M$ is called the equilibrium point of the mean field $\mathbf{x}(\cdot)$ if it satisfies $\mathbf{l}(\mathbf{P}) = \mathbf{0}$, i.e., $\mathbf{x}(t,\mathbf{P}) = \mathbf{P}$ for all $t \geq 0$. Hence, from (2.66) we have that for all $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$

$$P_{k+1,j} - P_{k+2,j} = \Delta_j \left((P_{k,j})^{d_j} - (P_{k+1,j})^{d_j} \right) \prod_{i=1}^{j-1} (P_{k,i})^{d_i} \prod_{i=j+1}^{M} (P_{k+1,i})^{d_i}, \tag{2.85}$$

where $\Delta_j = \frac{\lambda}{\mu\gamma_j C_j}$ for each $j \in \mathcal{J}$. Note that by definition we have $P_{0,j} = 1$ for all $j \in \mathcal{J}$. In the next proposition, we show that for a fixed $j \in \mathcal{J}$ the components $P_{k,j}$, $k \in \mathbb{Z}_+$ decrease with k super-exponentially, in the sense of the following definition.

Definition 2.6.1. *A real sequence $\{z_n\}_{n\in\mathbb{Z}_+}$ is said to decrease super-exponentially if and only if there exist positive constants L, $\omega < 1$, $\theta > 1$, and κ such that $z_n \leq \kappa\omega^{\theta^n}$ for all $n \geq L$.*

Clearly, if a sequence $\{z_n\}_{n\in\mathbb{Z}_+}$ decays super-exponentially, then it is summable, i.e., $\sum_{n=0}^{\infty} z_n < \infty$.

Proposition 2.6.2. *Let $\mathbf{P}$ be an equilibrium point of the mean field $\mathbf{x}(\cdot)$. Assume that for each $j \in \mathcal{J}$, $P_{k,j} \downarrow 0$ as $k \to \infty$. Then the following equations must hold*

$$\sum_{j\in\mathcal{J}} \frac{P_{l+1,j}}{\Delta_j} = \prod_{j\in\mathcal{J}} (P_{l,j})^{d_j} . \tag{2.86}$$

Further, for each $j \in \mathcal{J}$, the sequence $\{P_{k,j}, k \in \mathbb{Z}_+\}$ decreases super exponentially. In particular, under the assumption of the proposition, $\{P_{k,j}, k \in \mathbb{Z}_+\}$ is a summable sequence.

Proof. For a fix j adding (2.85) for all $k \geq l$ and using the fact $\lim_{k\to\infty} P_{k,j} = 0$ we obtain

$$P_{l+1,j} = \Delta_j \sum_{k\geq l} \left[\prod_{i=1}^{j} (P_{k,i})^{d_i} \prod_{i=j+1}^{M} (P_{k+1,i})^{d_i} - \prod_{i=1}^{j-1} (P_{k,i})^{d_i} \prod_{i=j}^{M} (P_{k+1,i})^{d_i} \right] \tag{2.87}$$

Now, multiplying both sides of the above equation by $\frac{1}{\Delta_j}$ and adding the resulting equation over all $j \in \mathcal{J}$ (and using $\lim_{k\to\infty} P_{k,j} = 0$) we obtain (2.86). From (2.86) we obtain $\frac{P_{k+1,j}}{\Delta_j} \leq \prod_{j\in\mathcal{J}} (P_{k,j})^{d_j} \leq \left(\hat{P}_k\right)^d$, where $\hat{P}_k = \max_{1\leq j\leq M} P_{k,j}$ and $d = \sum_{j\in\mathcal{J}} d_j$. Thus, we have $P_{k+1,j} \leq \delta\hat{P}_k$, where $\delta = \left(\hat{P}_k\right)^{d-1} \max_{1\leq j\leq M}(\Delta_j)$. Since by hypothesis, for each j, $P_{k,j} \to 0$ as $k \to \infty$, one can choose k sufficiently large such that $\delta < 1$. Hence, we have $(\max_{1\leq j\leq M} P_{k+1,j}) \leq \delta\hat{P}_k$. Similarly we have, $(\max_{1\leq j\leq M} P_{k+n,j}) \leq \delta^{\frac{d^n-1}{d-1}}\hat{P}_k$. This proves that the sequence $\{P_{k,j}, k \in \mathbb{Z}_+\}$ decreases doubly exponentially for each j. □

The following theorem guarantees the existence of an equilibrium point $\mathbf{P}$ of the process $\mathbf{x}(\cdot)$.

Theorem 2.6.3. *If $\lambda \in \Lambda$, then there exists an equilibrium point $\mathbf{P} \in \mathcal{U}^M$ of the process $\mathbf{x}(\cdot)$.*

Proof. The proof is given in Appendix B. □

The next theorem shows that $\mathbf{P}$ is unique and globally asymptotically stable.

Theorem 2.6.4. *If $\lambda \in \Lambda$, then*

$$\lim_{t\to\infty} \mathbf{x}(t, \mathbf{u}) = \mathbf{P} \in \mathcal{U}^M \text{ for all } \mathbf{u} \in \mathcal{U}^M, \tag{2.88}$$

Furthermore, $\mathbf{P}$ is the only equilibrium point of $\mathbf{x}(\cdot)$ in the space $\mathcal{U}^M$.

To prove Theorem 2.6.4, we first state the following lemma. We will write $\mathbf{u} \le \mathbf{u}'$ to mean that $u_{n,j} \le u'_{n,j}$ holds for all $n \in \mathbb{Z}_+$ and $j \in \mathcal{J}$.

Lemma 2.6.3. *If $\mathbf{u} \le \mathbf{u}'$ holds, for $\mathbf{u}, \mathbf{u}' \in \bar{\mathcal{U}}^M$, then $\mathbf{x}(t, \mathbf{u}) \le \mathbf{x}(t, \mathbf{u}')$ holds for all $t \ge 0$.*

Proof. We observe from (2.66) that for each $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$, $l_{k,j}(\mathbf{x})$ is a non-decreasing in $x_{n,i}$ for all $(n, i) \ne (k, j)$. In other words, $dx_{k,j}(t)/dt$ is quasi-monotone for each $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$. Therefore, the statement of the lemma follows from [[38], p. 70-74]. □

We define $z_{n,j}(t, \mathbf{u}) = \sum_{k\ge n} x_{k,j}(t, \mathbf{u})$ and $z_n(t, \mathbf{u}) = \sum_{j\in\mathcal{J}} \gamma_j z_{n,j}(t, \mathbf{u})$ for each $n \ge 1$ and $j \in \mathcal{J}$. Further, $z_{n,j}(\mathbf{u}) = \sum_{k\ge n} u_{k,j}$ and $z_n(\mathbf{u}) = \sum_{j\in\mathcal{J}} \gamma_j z_{n,j}(\mathbf{u})$ for each $n \ge 1$ and $j \in \mathcal{J}$.

Lemma 2.6.4. *If $\mathbf{u} \in \mathcal{U}^M$, then $\mathbf{x}(t, \mathbf{u}) \in \mathcal{U}^M$ for all $t \ge 0$ and*

$$\frac{dz_n(t, \mathbf{u})}{dt} = \lambda \left(\prod_{j=1}^{M} (x_{n-1,j}(t, \mathbf{u}))^{d_j} - \sum_{j=1}^{M} \frac{x_{n,j}(t, \mathbf{u})}{\Delta_j} \right) \text{ for all } n \ge 1. \tag{2.89}$$

In particular,

$$\frac{dz_1(t, \mathbf{u})}{dt} = \lambda \left(1 - \sum_{j=1}^{M} \frac{x_{1,j}(t, \mathbf{u})}{\Delta_j} \right) \tag{2.90}$$

Proof. The result directly follows from (2.66) by first summing the right hand side of (2.66) over all $k \ge n$ and then over all $j \in \mathcal{J}$. □

Using the above results we now prove Theorem 2.6.4 as follows.

Proof of Theorem 2.6.4: Clearly, Lemma 2.6.3 implies the following

$$\mathbf{x}(t, \min(\mathbf{u}, \mathbf{P})) \leq \mathbf{x}(t, \mathbf{u}) \leq \mathbf{x}(t, \max(\mathbf{u}, \mathbf{P})) \tag{2.91}$$

Hence, to prove (2.88), it is sufficient to show that the convergence holds for $\mathbf{u} \geq \mathbf{P}$ and for $\mathbf{u} \leq \mathbf{P}$.

We first need to check that for each such $\mathbf{u}$, the quantity $z_1(t, \mathbf{u})$ (and hence also $z_n(t, \mathbf{u})$ for $n > 1$) is bounded uniformly in t. If $\mathbf{u} \leq \mathbf{P}$, then by Lemma 2.6.3 we have $\mathbf{x}(t, \mathbf{u}) \leq \mathbf{x}(t, \mathbf{P}) = \mathbf{P}$ for all $t \geq 0$. Hence, $z_1(t, \mathbf{u}) \leq z_1(\mathbf{P})$.

On the other hand, if $\mathbf{u} \geq \mathbf{P}$, then by Lemma 2.6.3, we have $\mathbf{x}(t, \mathbf{u}) \geq \mathbf{x}(t, \mathbf{P}) = \mathbf{P}$. Hence,

$$\sum_{j=1}^{M} \frac{x_{1,j}(t, \mathbf{u})}{\Delta_j} \geq \sum_{j=1}^{M} \frac{P_{1,j}}{\Delta_j} = 1. \tag{2.92}$$

Thus, from (2.90) we have $\frac{dz_1(t,\mathbf{u})}{dt} \leq 0$. This implies that $0 \leq z_1(t, \mathbf{u}) \leq z_1(\mathbf{u})$ for all $t \geq 0$.

Since the derivative of $x_{n,j}(t)$ is bounded for all $j \in \mathcal{J}$, the convergence $\mathbf{x}(t, \mathbf{u}) \to \mathbf{P}$ will follow from

$$\int_0^\infty (x_{n,j}(t, \mathbf{u}) - P_{n,j})\, dt < \infty, \;\; j \in \mathcal{J}, n \geq 1 \tag{2.93}$$

in the case $\mathbf{u} \geq \mathbf{P}$, and from

$$\int_0^\infty (P_{n,j} - x_{n,j}(t, \mathbf{u}))\, dt < \infty, \;\; j \in \mathcal{J}, n \geq 1 \tag{2.94}$$

in the case $\mathbf{u} \leq \mathbf{P}$. Both the bounds can be shown similarly. We discuss the proof of (2.93).

To prove (2.93) it is sufficient to show that

$$\int_0^\infty \sum_{j=1}^{M} \frac{(x_{n,j}(t, \mathbf{u}) - P_{n,j})}{\Delta_j} dt < \infty, \tag{2.95}$$

for all $n \geq 1$. We show this using induction starting with $n = 1$. Using (2.90), we have

$$
\begin{aligned}
\int_0^\tau \sum_{j=1}^M \frac{(x_{1,j}(t,\mathbf{u}) - P_{1,j})}{\Delta_j} dt &= \left(\int_0^\tau \sum_{j=1}^M \frac{x_{1,j}(t,\mathbf{u})}{\Delta_j} - 1 \right) dt \\
&= -\frac{1}{\lambda} \int_0^\tau \frac{dz_1(t,\mathbf{u})}{dt} dt \\
&= \frac{1}{\lambda}(z_1(\mathbf{u}) - z_1(\tau,\mathbf{u})).
\end{aligned}
$$

Since the right hand side is bounded by a constant for all τ, the integral on the left hand side must converge as $\tau \to \infty$.

Now assume that (2.93) holds for all $n \leq L-1$. We have from (2.89) and (2.86)

$$
\begin{aligned}
z_L(0,\mathbf{u}) - z_L(\tau,\mathbf{u}) &= -\int_0^\tau \frac{dz_L(t,\mathbf{u})}{dt} dt \\
&= \lambda \int_0^\tau \left(\sum_{j=1}^M \frac{x_{L,j}(t,\mathbf{u})}{\Delta_j} - \prod_{j=1}^M (x_{L-1,j}(t,\mathbf{u}))^{d_j} \right) dt \\
&= \lambda \int_0^\tau \sum_{j=1}^M \frac{(x_{L,j}(t,\mathbf{u}) - P_{L,j})}{\Delta_j} dt \\
&\qquad + \lambda \int_0^\tau \left(\sum_{j=1}^M \frac{P_{L,j}}{\Delta_j} - \prod_{j=1}^M (x_{L-1,j}(t,\mathbf{u}))^{d_j} \right) dt \\
&= \lambda \int_0^\tau \sum_{j=1}^M \frac{(x_{L,j}(t,\mathbf{u}) - P_{L,j})}{\Delta_j} dt \\
&\qquad - \lambda \int_0^\tau \left(\prod_{j=1}^M (x_{L-1,j}(t,\mathbf{u}))^{d_j} - \prod_{j=1}^M (P_{L-1,j})^{d_j} \right) dt
\end{aligned}
$$

By the induction hypothesis, the last integral on the right hand side converges as $\tau \to \infty$. The left hand side also is uniformly bounded. Hence, the first integral on the left hand side also must converge as required. □

Remark 11. A distribution $\boldsymbol{\pi} \in \mathcal{P}(\bar{\mathcal{U}}^M)$ is called an invariant distribution of the map $\mathbf{u} \to \mathbf{x}(t,\mathbf{u})$ if for all continuous (and hence bounded) functions $\varphi : \bar{\mathcal{U}}^M \to \mathbb{R}$ and all $t \geq 0$ we have

$$\int \varphi(\mathbf{x}(t,\mathbf{u}))d\boldsymbol{\pi}(\mathbf{u}) = \int \varphi(\mathbf{u})d\boldsymbol{\pi}(\mathbf{u}). \tag{2.96}$$

Hence, by definition $\delta_{\mathbf{P}}$, the Dirac measure concentrated at the equilibrium point $\mathbf{P}$, is an invariant distribution for the map $\mathbf{u} \to \mathbf{x}(t,\mathbf{u})$. Conversely, suppose that $\boldsymbol{\pi}$ is an invariant distribution of the map $\mathbf{u} \to \mathbf{x}(t,\mathbf{u})$ with $\boldsymbol{\pi}(\mathcal{U}^M) = 1$. Let $\mathbf{u}$ be chosen according to $\boldsymbol{\pi}$. Theorem 2.6.4 implies that $\mathbf{x}(t,\mathbf{u}) \to \mathbf{P}$ as $t \to \infty$ for all $\mathbf{u} \in \mathcal{U}^M$. This implies that for all continuous (hence bounded)functions $\varphi : \bar{\mathcal{U}}^M \to \mathbb{R}$ and all $\mathbf{u} \in \mathcal{U}^M$, $\varphi(\mathbf{x}(t,\mathbf{u})) \to \varphi(\mathbf{P})$ as $t \to \infty$. By dominated convergence theorem this implies that $\lim_{t\to\infty} \int \varphi(\mathbf{x}(t,\mathbf{u}))d\boldsymbol{\pi}(\mathbf{u}) = \int(\lim_{t\to\infty} \varphi(\mathbf{x}(t,\mathbf{u})))d\boldsymbol{\pi}(\mathbf{u}) = \varphi(\mathbf{P})$. But since $\boldsymbol{\pi}$ is assumed to be an invariant map of $\mathbf{u} \to \mathbf{x}(t,\mathbf{u})$ we also have

$$\lim_{t\to\infty} \int \varphi(\mathbf{x}(t,\mathbf{u}))d\boldsymbol{\pi}(\mathbf{u}) = \int \varphi(\mathbf{u})d\boldsymbol{\pi}(\mathbf{u}) \tag{2.97}$$

Thus, we have $\int \varphi(\mathbf{u})d\boldsymbol{\pi}(\mathbf{u}) = \varphi(\mathbf{P})$ for all continuous maps φ. This implies $\boldsymbol{\pi} = \delta_{\mathbf{P}}$. Hence, a probability measure $\boldsymbol{\pi} \in \mathcal{P}(\bar{\mathcal{U}}^M)$ with $\boldsymbol{\pi}(\mathcal{U}^M) = 1$ is an invariant measure of the map $\mathbf{u} \to \mathbf{x}(t,\mathbf{u})$ if and only if $\boldsymbol{\pi} = \delta_{\mathbf{P}}$.

We now show that, if $\lambda \in \Lambda$, the stationary distribution of the process $\mathbf{x}^{(N)}(\cdot)$ (which exists and is unique due to stability) converges weakly to the Dirac measure concentrated at the unique equilibrium point of the mean field. Let $\boldsymbol{\pi}_N \in \mathcal{P}(\bar{\mathcal{U}}^M)$ denote the stationary distribution of the process $\mathbf{x}^{(N)}(\cdot)$. Furthermore, let $\mathbf{x}^{(N)}(\infty) = \lim_{t\to\infty} \mathbf{x}^{(N)}(t)$ denote the equilibrium state of the finite system distributed according to $\boldsymbol{\pi}_N$.

Theorem 2.6.5. *If $\lambda \in \Lambda$, then $\boldsymbol{\pi}_N \Rightarrow \delta_{\mathbf{P}}$ or equivalently $\mathbf{x}^{(N)}(\infty) \Rightarrow \mathbf{P}$.*

Proof. We recall that the space $\bar{\mathcal{U}}^M$ is compact under the metric ω. Hence, the sequence of probability measures $(\boldsymbol{\pi}_N)_N \in \mathcal{P}(\bar{\mathcal{U}}^M)$ is tight (See definition C.1.3). Thus Theorem C.1.2, implies that $(\boldsymbol{\pi}_N)_N$ is relatively compact and thus has limit points. In order to prove the theorem, we now need to show that all limit points coincide with $\delta_{\mathbf{P}}$.

Suppose that a subsequence $(\boldsymbol{\pi}_{N_k})_k$ of the sequence $(\boldsymbol{\pi}_N)_N$ converges to the limiting distribution $\boldsymbol{\pi}$. Further, for each k let the process $\mathbf{x}^{(N_k)}(\cdot)$, start with initial distribution $\boldsymbol{\pi}_{N_k}$ and let the mean field process $\mathbf{x}(\cdot)$ start with distribution $\boldsymbol{\pi}$. By the convergence of the operator semigroups established in Theorem 2.6.2 and Theorem C.2.2, it follows that $\mathbf{x}^{(N_k)}(t) \Rightarrow \mathbf{x}(t)$ for all $t \geq 0$. Since, $\boldsymbol{\pi}_{N_k}$ is the stationary distribution of the process $\mathbf{x}^{(N_k)}(\cdot)$, the distribution of $\mathbf{x}^{(N_k)}(t)$ is $\boldsymbol{\pi}_{N_k}$ for all $t \geq 0$. Hence, $\boldsymbol{\pi}_{N_k} \Rightarrow \boldsymbol{\pi}$ and $\mathbf{x}^{(N_k)}(t) \Rightarrow \mathbf{x}(t)$ together imply that the distribution of $\mathbf{x}(t)$ must be $\boldsymbol{\pi}$ for all $t \geq 0$. This implies

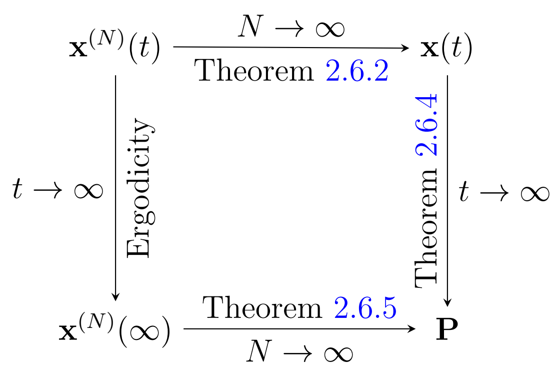

Figure 2.2: Commutativity of limits

that $\boldsymbol{\pi}$ is an invariant distribution of the map $\mathbf{u} \mapsto \mathbf{x}(t, \mathbf{u})$. If we can show that $\boldsymbol{\pi}$ satisfies $\boldsymbol{\pi}(\mathcal{U}^M) = 1$, then by Remark 11 we can conclude $\boldsymbol{\pi} = \delta_{\mathbf{P}}$ which proves the theorem. To prove that $\boldsymbol{\pi}$ is concentrated on $\mathcal{U}^M$ it is sufficient to show that $\mathbb{E}_{\pi}\left[\sum_{n\geq 1} u_{n,j}\right] < \infty$ for all $j \in \mathcal{J}$. The coupling described in Remark 10 implies that $\mathbb{E}_{\pi_N}\left[\sum_{n\geq 1} u_{n,j}\right] \leq \frac{\rho}{1-\rho}$, where $\rho = \frac{\lambda}{\mu \sum_{j\in\mathcal{J}} \gamma_j C_j} < 1$. Hence, $\mathbb{E}_{\pi}\left[\sum_{n\geq 1} u_{n,j}\right] = \lim_{N\to\infty} \mathbb{E}_{\pi_N}\left[\sum_{n\geq 1} u_{n,j}\right] \leq \frac{\rho}{1-\rho}$. This completes the proof. □

We have therefore established that the interchange of limits indicated in Figure 2.2 holds.

2.6.1 Propagation of chaos

So far we have considered the convergence of the process $\mathbf{x}^{(N)}(\cdot)$ which describes the evolution of the entire system. In this subsection, we focus on a given finite set of servers in the system. We show that as the system size grows the servers in the set become mutually independent. Such independence is shown to hold at any finite time and at the equilibrium, provided that the initial server occupancies satisfy certain assumptions. We also find the stationary distribution of occupancy of each server in the large system limit using the independence property. It is shown that the stationary distribution is given by the unique equilibrium point $\mathbf{P}$ of the mean field $\mathbf{x}(\cdot)$. The independence property considered is known as the *propagation of chaos* [29, 11] or *asymptotic independence property* [34, 32] in the literature.

To formally state the results, we first introduce the following notations.

- Let $q_{k,j}^{(N)}(t)$ and $q_{k,j}^{(N)}(\infty)$, for $k \in \{1, 2, \dots, N\gamma_j\}$ and $j \in \mathcal{J}$, denote the occupancy of the k^{th} server of type j at time $t \geq 0$ and at equilibrium, respectively.
- For $j \in \mathcal{J}$ and $n \in \mathbb{Z}_+$, let $\bar{x}_{n,j}^{(N)}(t) = x_{n,j}^{(N)}(t) - x_{n+1,j}^{(N)}(t)$ denote the fraction of type j servers having occupancy exactly n at time $t \geq 0$. Define the process $\bar{\mathbf{x}}^{(N)}(\cdot) = \left(\bar{x}_{n,j}^{(N)}(\cdot), n \in \mathbb{Z}_+, j \in \mathcal{J}\right)$. For each j, the vector $\bar{\mathbf{x}}_j^{(N)}(t) = \left(\bar{x}_{n,j}^{(N)}(t), n \in \mathbb{Z}_+\right)$, therefore, denotes the empirical distribution of occupancies of type j servers in the finite system. By $\bar{\mathbf{x}}_j^{(N)}(\infty)$ we shall denote the empirical distribution of occupancies of type j servers at equilibrium.
- Let the process $\bar{\mathbf{x}}(\cdot) = (\bar{x}_{n,j}(\cdot), j \in \mathcal{J}, n \in \mathbb{Z}_+)$ be defined as $\bar{x}_{n,j}(t) = x_{n,j}(t) - x_{n+1,j}(t)$, for $t \geq 0$. We define $\bar{x}_{n,j}(\infty) = P_{n,j} - P_{n+1,j}$, where $\mathbf{P}$ is the unique equilibrium point of the process $\mathbf{x}(\cdot)$. For $t \in [0, \infty]$, we denote by $\bar{\mathbf{x}}_j(t)$ the distribution on $\mathbb{Z}_+$ given by $\bar{\mathbf{x}}_j(t) = (\bar{x}_{n,j}(t), n \in \mathbb{Z}_+)$.

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

Definition 2.6.2. Let $\left(q_{k,j}^{(N)}, 1 \leq k \leq N\gamma_j, 1 \leq j \leq M\right)$ denote a collection of N random variables classified into M different types, where the index j represents the type. The collection is called *intra-type exchangeable* if the joint law of the collection is invariant under permutation of indices of random variables belonging to the same type. Thus, the collection $\left(q_{k,j}^{(N)}, 1 \leq k \leq N\gamma_j, 1 \leq j \leq M\right)$ is intra-type exchangeable if for each $i \in \{1, 2, \dots, M\}$ and permutation σ_i of the numbers $\{1, 2, \dots, N\gamma_i\}$, we have

$$\text{Law}\left(q_{k,j}^{(N)}, 1 \leq k \leq N\gamma_j, 1 \leq j \leq M\right) = \text{Law}\left(q_{\sigma_i(1),i}^{(N)}, q_{\sigma_i(2),i}^{(N)}, \dots q_{\sigma_i(N\gamma_i),i}^{(N)}; q_{k,j}^{(N)}, 1 \leq k \leq N\gamma_j,\, j \in \mathcal{J}\backslash\{i\}\right) \quad (2.98)$$

Theorem 2.6.6. *For the model considered in this paper if* $(q_{k,j}^{(N)}(0), 1 \leq k \leq N\gamma_j, 1 \leq j \leq M)$ *is intra-class exchangeable and if* $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u} \in \mathcal{U}^M$ *as* $N \to \infty$, *then the following holds*

(i) For each $j \in \mathcal{J}$, $1 \leq k \leq N\gamma_j$, *and* $t \in [0, \infty]$, $q_{k,j}^{(N)}(t) \Rightarrow U_j(t)$ *as* $N \to \infty$, *where* $U_j(t)$ *is a random variable with distribution* $\bar{\mathbf{x}}_j(t)$.

(ii) *Fix positive integers* $1 \leq r_1 \leq N\gamma_1, 1 \leq r_2 \leq N\gamma_2, \ldots, 1 \leq r_M \leq N\gamma_M$. *Then for each* $t \in [0, \infty]$,

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

as $N \to \infty$, *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* $\bar{\mathbf{x}}_j(t)$ *for all* $1 \leq k \leq r_j$.

Proof. Note that the first part of Theorem 2.6.6 is a special case of the second part. Hence, it is sufficient to prove the second part. For simplicity of notations, we shall prove the second part for the $M = 2$ case. The proof readily extends to any $M \geq 2$.

Since under the type based scheme all the servers of the same type are statistically identical, the collection $(q_{k,j}^{(N)}(t), 1 \leq k \leq N\gamma_j, 1 \leq j \leq 2)$ is intra-type exchangeable at time $t \in [0, \infty]$ provided that the collection $(q_{k,j}^{(N)}(0), 1 \leq k \leq N\gamma_j, 1 \leq j \leq M)$ is also intra-type exchangeable. [7]

Now, given that $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \mathcal{U}^M$ as $N \to \infty$ we know from Theorem 2.6.2 and Theorem 2.6.5 that $\bar{\mathbf{x}}^{(N)}(t) \Rightarrow \bar{\mathbf{x}}(t)$ as $N \to \infty$ for all $t \in [0, \infty]$. Henceforth, we will omit the variables t in our calculations since they hold for all $t \in [0, \infty]$. To prove the second part of the theorem for $M = 2$ it is sufficient to show that:

$$\mathbb{E}\left[\prod_{k=1}^{r_1} \phi_k\left(q_{k,1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{k,2}^{(N)}\right)\right] \to \prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2 \rangle \text{ as } N \to \infty \tag{2.99}$$

for all bounded mappings $\phi_k : \mathbb{Z}_+ \to \mathbb{R}_+$ and $\psi_k : \mathbb{Z}_+ \to \mathbb{R}_+$. We have

$$\begin{aligned}
&\left|\mathbb{E}\left[\prod_{k=1}^{r_1} \phi_k\left(q_{k,1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{k,2}^{(N)}\right)\right] - \prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2 \rangle\right| \\
&\quad \leq \left|\mathbb{E}\left[\prod_{k=1}^{r_1} \phi_k\left(q_{k,1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{k,2}^{(N)}\right)\right] - \mathbb{E}\left[\prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2^{(N)} \rangle\right]\right| \\
&\qquad + \left|\mathbb{E}\left[\prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_1^{(N)} \rangle\right] - \prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1 \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2 \rangle\right|,
\end{aligned} \tag{2.100}$$

[7] We note that the collection $(q_{k,j}^{(N)}(\infty), 1 \leq k \leq N\gamma_j, 1 \leq j \leq M)$ is intra-type exchangeable since the stationary distribution can be symmetrized by using appropriate scaling factor.

where $\bar{\mathbf{x}}_j^{(N)} = (x_{n,j}, n \in \mathbb{Z}_+)$ is the random probability measure on $\mathbb{Z}_+$ induced by the process $\bar{\mathbf{x}}^{(N)}(\cdot)$. We note that the second term on the right hand side of the above inequality vanishes as $N \to \infty$ because of the following facts: $\bar{\mathbf{x}}_j^{(N)} \Rightarrow \bar{\mathbf{x}}_j$ as $N \to \infty$ for $j = 1, 2$; $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$ are deterministic; $\bar{\mathbf{x}}_j^{(N)}$ is a bounded random vector for each $j = 1, 2$. 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

$$\mathbb{E}\left[\prod_{k=1}^{r_1} \phi_k\left(q_{k,1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{k,2}^{(N)}\right)\right] = \frac{1}{(N\gamma_1)_{r_1}(N\gamma_2)_{r_2}} \times \mathbb{E}\left[\sum_{\substack{\sigma \in P(r_1, N\gamma_1) \\ \sigma' \in P(r_2, N\gamma_2)}} \prod_{k=1}^{r_1} \phi_k\left(q_{\sigma(k),1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{\sigma'(k),2}^{(N)}\right)\right] \tag{2.101}$$

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 $\bar{\mathbf{x}}_j^{(N)}$ we have

$$\begin{aligned}\mathbb{E}&\left[\prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2^{(N)} \rangle\right] \\ &= \mathbb{E}\left[\left(\prod_{k=1}^{r_1} \frac{1}{N\gamma_1} \sum_{l=1}^{N\gamma_1} \phi_k\left(q_{l,1}^{(N)}\right)\right) \left(\prod_{k=1}^{r_2} \frac{1}{N\gamma_2} \sum_{l=1}^{N\gamma_2} \psi_k\left(q_{l,2}^{(N)}\right)\right)\right]\end{aligned} \tag{2.102}$$

Hence, the first term on the right hand side of (2.100) can be bounded as follows

$$\begin{aligned}&\left|\mathbb{E}\left[\prod_{k=1}^{r_1} \phi_k\left(q_{k,1}^{(N)}\right) \prod_{k=1}^{r_2} \psi_k\left(q_{2,k}^{(N)}\right)\right] - \mathbb{E}\left[\prod_{k=1}^{r_1} \langle \phi_k, \bar{\mathbf{x}}_1^{(N)} \rangle \prod_{k=1}^{r_2} \langle \psi_k, \bar{\mathbf{x}}_2^{(N)} \rangle\right]\right| \\ &\leq (N\gamma_1)_{r_1}(N\gamma_2)_{r_2} \left(\frac{1}{(N\gamma_1)_{r_1}(N\gamma_2)_{r_2}} - \frac{1}{(N\gamma_1)^{r_1}(N\gamma_2)^{r_2}}\right) B^{r_1+r_2} \\ &+ \left((N\gamma_1)^{r_1}(N\gamma_2)^{r_2} - (N\gamma_1)_{r_1}(N\gamma_2)_{r_2}\right) \frac{B^{r_1+r_2}}{(N\gamma_1)^{r_1}(N\gamma_2)^{r_2}} \\ &\leq 2B^{r_1+r_2}\left(1 - \frac{(N\gamma_1)_{r_1}(N\gamma_2)_{r_2}}{(N\gamma_1)^{r_1}(N\gamma_2)^{r_2}}\right) \to 0 \text{ as } N \to \infty\end{aligned}$$

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

Thus, the above proposition shows that in the limiting system server occupancies become independent of each other. It also shows that the stationary occupancy distribution of any given type j server in the limiting system is $\bar{\mathbf{x}}_j(\infty) = \{P_{n,j} - P_{n+1,j}, n \in \mathbb{Z}_+\}$. Hence, in the limiting system the stationary tail distribution of occupancy of a type $j \in \mathcal{J}$ server in the limiting system is given by $\{P_{k,j}, k \in \mathbb{Z}_+\}$. Using the independence of servers in the limiting system we conclude the following proposition.

Proposition 2.6.3. *In equilibrium, the arrival process of jobs at any given server in the limiting system is a state dependent Poisson process. Further, the arrival rate of jobs at a server of type $j \in \mathcal{J}$ when it has occupancy k in the equilibrium is given by*

$$\lambda_{k,j} = \frac{\lambda}{\gamma_j} \frac{(P_{k,j})^{d_j} - (P_{k+1,j})^{d_j}}{P_{k,j} - P_{k+1,j}} \prod_{i=1}^{j-1} (P_{k,i})^{d_i} \prod_{i=j+1}^{M} (P_{k+1,i})^{d_i} . \tag{2.103}$$

Proof. Consider a *tagged* type j server in the system and the arrivals that have the tagged server as one of its possible destinations. These arrivals constitute the *potential arrival process* at the tagged server. The probability that the tagged server is selected as a potential destination server for a new arrival is $\frac{\binom{N\gamma_j - 1}{d_j - 1}}{\binom{N\gamma_j}{d_j}} = \frac{d_j}{N\gamma_j}$. Thus, due to Poisson thinning, the potential arrival process to the tagged server is a Poisson process with rate $\frac{d_j}{N\gamma_j} \times N\lambda = \frac{d_j\lambda}{\gamma_j}$.

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 servers. However, as $N \to \infty$, due to the asymptotic independence property shown in Theorem 2.6.6 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 $N \to \infty$.

Consider the potential arrivals that occur to the tagged server when its occupancy is k. This arrival actually joins the tagged server with probability $\frac{1}{x+1}$ when x other servers among the d_j servers of type j have occupancy k, all the d_i servers of type $i < j$ have at least occupancy k, and all the d_i servers of type $i > j$ have at least occupancy $k+1$. Thus, the total arrival rate $\lambda_{k,j}$ can be computed as

$$\lambda_{k,j} = \frac{d_j \lambda}{\gamma_j} \sum_{x=0}^{d_j-1} \frac{1}{x+1} \binom{d_j-1}{x} (P_{k,j} - P_{k+1,j})^x (P_{k+1,j})^{d_j-1-x} \times \prod_{i=1}^{j-1} (P_{k,i})^{d_i} \prod_{i=j+1}^{M} (P_{k+1,i})^{d_i}, \quad (2.104)$$

which simplifies to (2.103). □

Thus, the above proposition shows that at equilibrium the arrival rate at a given server in the limiting system is a state dependent Poisson process whose rates depend on the stationary tail probabilities $P_{k,j}$, $k \in \mathbb{Z}_+$ and $j \in \mathcal{J}$ through (2.103).

The stationary tail probabilities can, in turn, be expressed as functions of the arrival rates. Indeed, in equilibrium the global balance equations (which hold under state dependent Poisson arrivals due to Theorems 3.10 and 3.14 of [39]) yield

$$(P_{k,j} - P_{k+1,j})\lambda_{k,j} = (P_{k+1,j} - P_{k+2,j})\mu C_j, \text{ for } j \in \mathcal{J}, k \in \mathbb{Z}_+. \quad (2.105)$$

We note that the above equation reduces to (2.85) if we replace the arrival rates $\lambda_{k,j}$ by the RHS of (2.103) The equilibrium point $\mathbf{P}$ is, therefore, the unique fixed point of the mapping $\Theta : \mathcal{U}^M \to \mathcal{U}^M$ defined as $\Theta(\mathbf{P}) = F(G(\mathbf{P}))$, where $G(\cdot)$ denotes the mapping from $\mathcal{U}^M$ to the space of possible arrival rates (defined by (2.103)) and $F(\cdot)$ denotes the mapping from the space of possible arrival rates to the space $\mathcal{U}^M$ (defined by (2.105)). Thus, the equilibrium point $\mathbf{P}$ can be computed using the fixed point iterations (i.e., by repeatedly applying the mapping $\Theta(\cdot)$ to some arbitrary point $\mathbf{Q} \in \mathcal{U}^M$.)

Remark 12. All the results discussed above have been obtained assuming exponential job length distributions. If the independence of servers shown in Theorem 2.6.6 holds for all job length distributions, then Proposition 2.6.3 continues to hold irrespective of the job length distribution. This implies that (2.105) holds. Since the servers in the system are processor sharing servers and (2.105) represents detailed balance, Theorem 1 of [40] implies that that the stationary distribution of each server in the limiting system is *insensitive* to job length distributions. Hence, under the assumption of asymptotic independence of servers for general job length distributions, the stationary distribution of server occupancies in the limiting system becomes insensitive to the job length distribution type and only depends on its mean. We refer to this as the *asymptotic insensitivity* property. The proof of asymptotic insensitivity requires proving asymptotic independence for general job length distributions, which remains as an open problem. Asymptotic independence was conjectured in [34] for homogeneous systems with any local service disciple (where the rate at which jobs are processed depends only on the current jobs that are being processed) and general job length distributions. Our numerical results, presented in Section 2.7, supporting asymptotic insensitivity, suggest that such asymptotic independence should also hold for heterogeneous systems considered in this chapter.

Remark 13. As in the case of the SQ(d) scheme with heterogeneous servers, it is easily seen (by applying Little's law) that the mean response time of jobs for the type based scheme in the heterogeneous case can be expressed as a function of the stationary tail probabilities $P_{k,j}$, $k \in \mathbb{Z}_+$, $j \in \mathcal{J}$ and is given by

$$\bar{T}^{\text{tb}}(\lambda) = \frac{1}{\lambda} \sum_{j=1}^{M} \gamma_j \sum_{k=1}^{\infty} P_{k,j} \tag{2.106}$$

Thus, the mean response time of jobs can be computed by first computing the stationary point $\mathbf{P}$ from equations (2.103) and (2.105) (using the fixed point method discussed above) and then using (2.106).

2.7 Numerical Results

In this section, we first investigate the accuracy of the mean field analysis of the type-based scheme in predicting the performance of the scheme for large but finite systems. We then numerically compare the mean response time of jobs under the different job assignment schemes discussed in this chapter. Finally, numerical evidence to support asymptotic insensitivity is also provided. All simulation results, presented in this section, are obtained by averaging 10,000 independent runs.

To investigate the accuracy of the mean field analysis of the type-based scheme, we compare the mean response time of jobs computed from (2.106) with that obtained by simulating the finite system for different values of N and d, where $d_j = d$ for all $j \in \mathcal{J}$. We choose the following parameter setting: $M = 2$, $\gamma_1 = \gamma_2 = 0.5$, $\mu = 1$, $C_1 = 2/3$, $C_2 = 4/3$. For the above parameter setting the maximal stability region of the system is given by $\Lambda = \{\lambda : 0 \leq \lambda < 1\}$. We choose $\lambda = 0.8$, which lies in the stability region. The results are shown in Table 2.1. As expected, the difference between the asymptotic results and the corresponding simulation results decreases with the increase in N. We also observe that for the same value of N, increasing d, increases the percentage of error between the simulation results and the results obtained from the mean field limit. This is because for finite N increasing d increases the correlation between the servers. This acts in opposition to the independence of servers in the limiting system. From the results it is clear that the mean field analysis quite accurately captures the behavior of finite systems under the type-based scheme.

Table 2.1: Accuracy of the mean field analysis of the type-based scheme

d	Asymptotic	$N = 20$	$N = 50$	$N = 100$	$N = 200$
2	1.3687	1.4695	1.3960	1.3720	1.3689
4	1.0960	1.2319	1.1492	1.1211	1.1055
6	1.0123	1.1595	1.0699	1.0396	1.0281
8	0.9732	1.1216	1.0328	1.0007	0.9847
10	0.9539	1.1064	1.0083	0.9788	0.9646

We now compare the mean response time of jobs under the different schemes discussed in this chapter. We take the following parameter values: $M = 2$, $\gamma_1 = \gamma_2 = 0.5$, $\mu = 1$, $C_1 = 2/3$, $C_2 = 4/3$, $d = 2$, $N = 200$. For these parameter values we have $\Lambda_\infty = \Lambda = \{\lambda : 0 \leq \lambda < 1\}$. Hence, in this setting, the SQ(d) scheme achieves the maximal stability region. We note that this is in accordance with Proposition 2.4.1 since the condition $d \geq \lceil C_2/C_1 \rceil$ is satisfied in this case. In Figure 2.3, we plot the mean sojourn time of jobs in the system as a function of λ for the different schemes discussed in this chapter. The routing probabilities for the randomized static scheme and the hybrid SQ(2) scheme are chosen to be the optimal routing probabilities (obtained using (2.9) and (2.45), respectively). We have also plotted the mean sojourn time of jobs under the memory based scheme proposed in [31], in which the least loaded server among the d servers, sampled for an arriving job, is kept in the memory and used as a potential destination server for the next arriving job. From Figure 2.3, we observe that the mean sojourn time of jobs is highest under the randomized static scheme and lowest under the type-based scheme. The performance of the hybrid SQ(2) scheme is seen to be close that of the type-based scheme. Thus, we conclude that the proposed schemes significantly outperform other existing randomized schemes.

We now consider a second set of parameter values given by $M = 2$, $\gamma_1 = \gamma_2 = 0.5$, $\mu = 1$, $C_1 = 1/3$, $C_2 = 5/3$, $d = 2$, $N = 200$. For the above parameter values we have $\Lambda_\infty = \{\lambda : 0 \leq \lambda < 2/3\} \subset \Lambda = \{\lambda : 0 \leq \lambda < 1\}$. Hence, in this case, the stability region

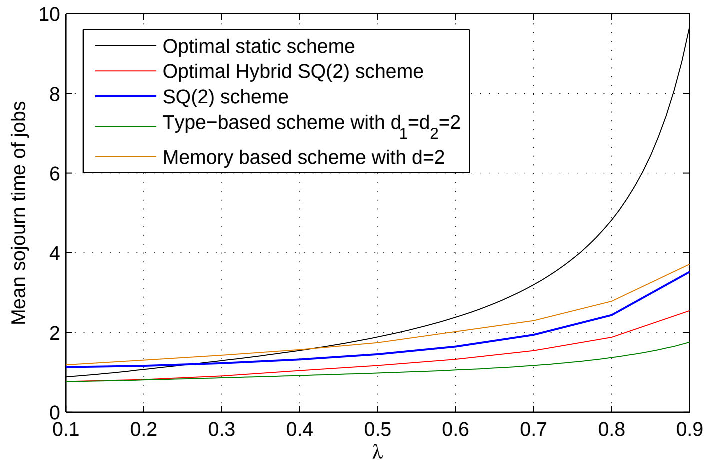

Figure 2.3: Mean sojourn time jobs as a function of λ for different schemes. Parameters: $M = 2$, $C_1 = 2/3$, $C_2 = 4/3$, and $\gamma_1 = \gamma_2 = 0.5$.

of the system under the SQ(d) scheme is smaller than the maximal stability region. In Figure 2.4, we plot the mean sojourn time of jobs in the system as a function of λ for the different schemes considered in this chapter. As in the previous case, the routing probabilities for the randomized static scheme and the hybrid SQ(2) scheme are chosen to be the optimal routing probabilities. We observe that in this case, the mean sojourn time of jobs under the randomized static scheme is lower than that under the SQ(2) scheme. This is expected since in this case the stability region of the system under the randomized static scheme is larger than that under the SQ(2) scheme. We observe that the hybrid SQ(2) scheme and the type-based scheme significantly outperform the randomized static scheme even in this case.

In Figures 2.3 and 2.4, the routing probabilities for the randomized static scheme and the hybrid SQ(2) scheme were chosen to be the optimal routing probabilities. Computing the optimal routing probabilities requires the knowledge of the the arrival rate λ, which is difficult to estimate. However, if we choose $p_i = \frac{\gamma_i C_i}{\sum_{j \in \mathcal{J}} \gamma_j C_j}$ for all $i \in \mathcal{J}$, then according to Remark 7 the system, operating under the hybrid SQ(d) scheme, is stable for any $\lambda \in \Lambda$. The same result also holds for the randomized static scheme. With the above choice of routing probabilities, in Figure 2.5 we now compare the mean sojourn time of jobs in the system under the randomized static scheme, the hybrid SQ(2) scheme and the type-based scheme. The parameters are chosen to be $M = 2$, $C_1 = 1/5$, $C_2 = 9/5$, $\gamma_1 = \gamma_2 = 0.5$. We observe that even in this case the performance of the type-based scheme is very close to

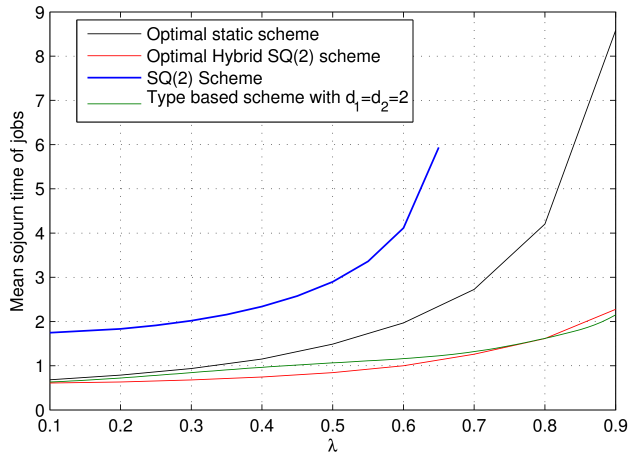

Figure 2.4: Mean sojourn time jobs as a function of λ for different schemes. Parameters: $M = 2$, $C_1 = 1/3$, $C_2 = 5/3$, and $\gamma_1 = \gamma_2 = 0.5$.

that of the hybrid SQ(2) scheme.

We now numerically investigate the behavior of the type based scheme under different job length distributions. In Table 2.2, mean sojourn time of jobs under the type-based scheme is shown as a function of λ, for the following distributions.

1. *Constant*: We consider the job length distribution having the cumulative distribution given by $F(x) = 0$ for $0 \leq x < 1$, and $F(x) = 1$, otherwise.
2. *Power law*: We consider the 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.

For both distributions we have $\mu = 1$. We choose the following parameter values $M = 2$, $C_1 = 4/3$, $C_2 = 2/3$, $N = 100$, $\gamma_1 = \gamma_2 = \frac{1}{2}$, and $d_1 = d_2 = 2$. We observe that there is no significant change in the mean sojourn time of jobs when the job length distribution type is changed. The results, therefore, numerically supports the asymptotic insensitivity property as discussed in Remark 12.

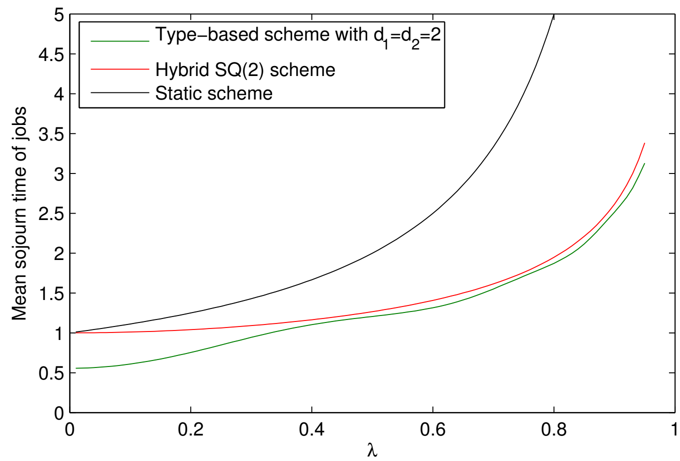

Figure 2.5: Mean sojourn time jobs as a function of λ for different schemes. Parameters: $M = 2$, $C_1 = 1/5$, $C_2 = 9/5$, $\gamma_1 = \gamma_2 = 0.5$, $N = 200$. The routing probabilities for the randomized static scheme and the hybrid SQ(2) scheme are not optimal.

2.8 Conclusion

In this chapter, we investigated the stability and performance of randomized dynamic job assignment schemes for heterogeneous processor sharing systems. We showed that uniform sampling of servers from the entire system may result in the reduction of the stability region of the system. A hybrid scheme in which biased sampling of different server types is combined with uniform sampling of servers within the same type was shown to achieve the maximal stability region but only with the knowledge of the system parameters. A type-based scheme in which servers of each type are sampled at each arrival instant of a new job was also considered. We saw that the type-based scheme achieves the maximal stability region without requiring the knowledge of the system parameters. The performance of the type-based scheme was characterized using mean field techniques. Numerical results were presented to show that the hybrid scheme and the type-based scheme have significantly lower mean response time of jobs than randomized static schemes. We also observed that under the proposed schemes the system is insensitive to the type of job length distribution in the large system limit. Some of our preliminary results on the stability of the SQ(d) scheme in the heterogeneous scenario and on the analysis of the hybrid SQ(d) scheme were presented in [41, 42, 43]. The detailed analysis of the type-based scheme appeared in [44].

Table 2.2: Asymptotic insensitivity of the type-based scheme

λ	Mean sojourn time (Mean Field)	Constant (Simulation)	Power Law (Simulation)
0.2	0.8076	0.8106	0.8098
0.3	0.8609	0.8642	0.8640
0.5	0.9809	0.9852	0.9840
0.7	1.1696	1.1759	1.1757
0.8	1.3687	1.3741	1.3740
0.9	1.7531	1.7641	1.7645

Chapter 3

Randomized Job Assignment in Heterogeneous Loss Systems

In this chapter, we consider randomized job assignment schemes for a system of parallel servers each having a finite amount of a resource. The incoming jobs are assumed to have specific resource requirements. Therefore, only a finite number of jobs can be processed simultaneously at each server. This is unlike the model in Chapter 2, where the incoming jobs had elastic resource requirements and therefore could be served by processor sharing servers. A job is accepted for processing at a given server only if the server has enough resource to process the job. Otherwise, the job is discarded or blocked. We consider randomized schemes to assign jobs to the 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 the maximum vacancy or maximum amount of unused resource among d servers, sampled uniformly at random. We show that the above scheme significantly reduces the average blocking probability of jobs as compared to randomized static schemes in which the incoming jobs are assigned to the servers independently of the states of the servers.

3.1 Introduction

Consider jobs with specific requirements of a resource, arriving at a multi-server system consisting of N parallel servers. Each server holds a finite amount of the required resource and therefore can process only a finite number of jobs simultaneously. We consider a *heterogeneous* system where different servers hold different amounts of the resource. Upon arrival, a job is routed/assigned to a server, where the job is either accepted or blocked depending on the availability of the resource requested by the job. If accepted, the processing of the job begins immediately at the server. Our aim is to design job assignment schemes that reduce the average blocking probability of the jobs while requiring the state information of a small number of servers at each job arrival instant.

Models, such as the one described above, arise frequently in the context of cloud computing systems that provide Infrastructure-as-a-service (IaaS) [45, 46]. 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 its users to choose from various classes of VM's (e.g. large, small etc.), differing in the amounts of resources they hold. We model situations [47], where the VM's differ only in one bottleneck resource (e.g. memory). In such situations, different classes of VM's correspond to different amounts of the same resource.

Each user, depending on its requirement, requests a VM of a specific class. 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 the requested resource. If accepted, the user holds the VM for the duration of its service after which it is released. To maintain a certain quality of service, the cloud service provider aims at reducing the average blocking probability of users which measures the fraction of time a user is denied access to its required resource.

The average blocking probability of jobs in the system can be reduced by suitably assigning the incoming jobs to the servers. Ideally, a job assignment scheme, which compares the states of all the servers at every arrival instant of a new job, has the potential to minimize the average blocking probability of the jobs. However, for large systems, such a scheme will involve high communication overhead between the job dispatcher and the servers. To reduce the overhead, we propose a randomized dynamic scheme in which only d ($\ll N$) servers are randomly sampled at each arrival instant and the arriving request is assigned to the server having the maximum available resource (vacancy) among the d sampled servers. We show that the above scheme results in a significant reduction in the average blocking probability of jobs as compared to randomized static job assignment schemes, where job assignments are made independently of the server states. Our results also show that at 'high' loads the the performance of the proposed scheme is nearly optimal in terms of reducing the average blocking probability of jobs.

3.1.1 Related literature

The dynamic routing scheme considered in this chapter is a loss model analog of the SQ(d) scheme considered in Chapter 2. A detailed discussion on the literature treating the SQ(d) scheme for first-come-first-serve servers is given in Section 2.1.1 of Chapter 2. Here we only discuss the works relevant to loss networks.

Turner [4, 48] studied the SQ(d) 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 shown explicitly.

Our work is closest to a recent work by Xie *et al.* [49], which analyzed an Erlang loss system with identical (homogeneous) servers under the SQ(d) policy using mean field techniques. In [49], the existence, uniqueness, and (global) asymptotic stability of the equilibrium point of the mean field were established for the homogeneous (servers) case with single class of customers. For the homogeneous (servers) case with multiple class of customers, the paper derived a recursive relationship among the tail probabilities of the number of occupied resource units for the limiting system. In this chapter, we generalize their results to the scenario where the servers have heterogeneous capacities. For this scenario, we establish independence of the servers in the limiting system through the milder requirement of *intra-type exchangeability* since the stronger requirement of exchangeability among all types of servers does not hold in the heterogeneous scenario. Such asymptotic independence of servers in the large system limit, also known as the *propagation of chaos* property, was studied earlier in the context of alternative routing by Graham and Méléard [50, 51] where the independence among servers was established on the path space of the processes of interest.

Mean field techniques have also been used in [52, 53] to study dynamic alternate routing policies for a fully connected, circuit switched network. The policies were analyzed in the limit as the number of links in the system increases to infinity. A simpler version of the problem, for a network without spatial features, was considered in [54]. A lattice caricature of the alternate routing problem was analyzed in [14]. In all the above works, it was found that, for certain ranges of the system parameters, multiple stable equilibrium points of the mean field exist. Hence, in such cases the finite system exhibits *metastability*, where the system fluctuates among multiple stable configurations. Metastability was also observed in [55] for an open network of loss servers where jobs of different classes move from one server to the other until they complete their service at all servers. However, for the system under consideration in this chapter, the mean field is shown to have a unique, asymptotically stable equilibrium point for all parameter settings.

3.1.2 Contributions

In this chapter, we consider the loss model analog of the SQ(d) scheme, where each incoming job is assigned to the server having the maximum vacancy among a set of d servers, sampled uniformly at random from the entire system. A job is discarded or blocked if none of the d servers, sampled at its arrival instant, has the required resource to process the job. We refer to this scheme as the MV(d) scheme, keeping in mind its similarity with the SQ(d) scheme discussed in the previous chapter.

The performance of the MV(d) scheme is analyzed for a system with heterogeneous servers and multiple job classes using mean field techniques. The mean field limit is shown to have a unique and asymptotically stable equilibrium point which characterizes the stationary distribution of the states of the servers in the limiting system. Using the concept of intra-type exchangeable random variables, introduced in Section 2.6.1 of Chapter 2, it is shown that any finite set of servers in the limiting system become independent.

Using the asymptotic independence of servers, a one dimensional recursive relationship between the stationary tail probabilities of server occupancies is obtained. This allows computation of the average blocking probability of jobs without having to compute the stationary distribution of the states of the servers. Using the recursive relationship, the rate of decay of stationary tail distribution of server occupancies is characterized. Numerical results are presented to show that the MV(d) scheme significantly reduces the average blocking probability of jobs as compared to the randomized static job assignment schemes. Furthermore, at 'high' loads the scheme is shown to be nearly optimal in terms reducing the average blocking probability of jobs.

3.1.3 Organization

The rest of the chapter is organized as follows. In Section 3.2, we introduce the detailed system model and the proposed job assignment scheme. In Section 3.3, we state our main results and discuss their implications. Section 3.4 presents the detailed proofs of the main results. Numerical results are provided in Section 3.5 to compare the proposed scheme with other existing schemes. Finally, the chapter is concluded in Section 3.6.

3.2 System model

We consider a system consisting of N parallel servers, where jobs or VM requests arrive and request necessary resource for processing. The servers are categorized into M different *types* based on the amounts of resource they hold or their *capacities*. Let $\mathcal{J} = \{1, 2, \ldots, M\}$ be the index set of server types. A server of type $j \in \mathcal{J}$ is assumed to hold C_j units of the resource. Without loss of generality, we assume that the capacities are ordered as follows:

$$C_1 \leq C_2 \leq \ldots \leq C_M. \tag{3.1}$$

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 $l \in \mathcal{L} = \{1, 2, \ldots, L\}$ VM requests require $A_l \geq 0$ units of resource from any given server in the system 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 $\underline{A} = (A_1, A_2, \ldots, A_L)$ the L-dimensional vector of resource requirements. We say that a server is in state $\underline{n} = (n_1, n_2, \ldots, n_L)$ 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 $\mathcal{S}_j = \left\{\underline{n} \in \mathbb{Z}_+^L : \underline{n} \cdot \underline{A} \leq C_j\right\}$, 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., $\mathcal{B}_j^{(l)} = \left\{\underline{n} \in \mathbb{Z}_+^L : C_j - A_l < \underline{n} \cdot \underline{A} \leq C_j\right\}$. Upon arrival, a job is routed to one of the N servers according to the following routing scheme:

The MV(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 job assigned to it only if the resource requested by the arriving job is available at the server. If accepted, processing of the job begins immediately. Otherwise, if the server is in a blocking state for the arriving request, then the request is discarded or blocked and lost. Clearly, a class-l VM request is blocked only when all the d potential destination servers have vacancies less than A_l. The service times of the 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.

3.3 Main results

In this section, we state the main results of this chapter and discuss their consequences. Our results are asymptotic in the sense that they are derived in the limit as the system size $N \to \infty$ keeping the proportions γ_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 having varied capacities. 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 3.2, let $P_{k,j}^{(N)}$ denote the stationary probability that a server of type $j \in \mathcal{J}$ has at least k units of occupied resource. Then $P_{k,j}^{(N)}$ converges to $P_{k,j}$ as $N \to \infty$, where $P_{k,j}$ is the solution of the following recursive relationship:

$$k\left(P_{k,j} - P_{k+1,j}\right) = \sum_{l=1}^{L} \frac{A_l \lambda_l}{\gamma_j} \left[\left(\sum_{i=1}^{j} \gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j+1}^{M} \gamma_i P_{k-A_l+C_i-C_j+1,i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j}^{M} \gamma_i P_{k-A_l+C_i-C_j+1,i} \right)^d \right], \quad (3.2)$$

where $1 \le k \le C_j$, $P_{k,j} = 1$ for $k \le 0$, and $P_{C_j+1,j} = 0$ for all $j \in \mathcal{J}$. Furthermore, in the limit as $N \to \infty$ the servers become mutually independent.

Remark 14. Using the independence of servers stated above and the probabilities $P_{k,j}$ found by solving (3.2) the blocking probability $P_{\text{blocking}}^{(l)}$ of class-l requests in the limiting system can be computed as follows: A class-l request is blocked at a server with capacity C_j if the number of units of available resource is less than A_l. The stationary probability that a server of type j has less than A_l units of available resource is $P_{C_j-A_l+1,j}$. The probability with which a type-j server is sampled at the arrival instant of a job γ_j (due to uniform sampling). Thus the total probability that a randomly sampled server is in a 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 a class-l arrival is blocked is given by $P^{(l)}_{\text{blocking}} = \left(\sum_{j\in\mathcal{J}}\gamma_j P_{C_j-A_l+1,j}\right)^d$. Hence, the one-dimensional recursive relation (3.2) allows one to compute the average blocking probability of jobs in the limiting system using the stationary probabilities of the total number of occupied resource units. The stationary probabilities of the total number of occupied resource units at a server are much simpler to compute than the stationary probabilities of the server states since the state of each server lies in a multi-dimensional state space.

A lower bound on the average blocking probability: Under an arbitrary job assignment policy, the average blocking probability of jobs can be lower bounded as follows. For an arbitrary job assignment scheme, let the average blocking probability of class-l requests be denoted by $P^{(l)}_{\text{blocking}}$. Hence, by Little's law, the average number of class-l requests in the system is given by $(1-P^{(l)}_{\text{blocking}})N\lambda_l$. Now let the maximum number of class-l requests that a server of type j can process be $B^{(l)}_j = \max_{\underline{n}\in\mathcal{S}_j}(n_l)$. Then the average number of class-l requests in the entire system is upper bounded by $N\sum_{j\in\mathcal{J}}\gamma_j B^{(l)}_j$. We therefore have

$$(1-P^{(l)}_{\text{blocking}})N\lambda_l \leq N\sum_{j\in\mathcal{J}}\gamma_j B^{(l)}_j, \tag{3.3}$$

from which we obtain the following lower bound on $P^{(l)}_{\text{blocking}}$.

$$P^{(l)}_{\text{blocking}} \geq \left(1-\frac{\sum_{j\in\mathcal{J}}\gamma_j B^{(l)}_j}{\lambda_l}\right)_+ = \left(1-\frac{\lambda^{(l)}_{\text{crit}}}{\lambda_l}\right)_+, \tag{3.4}$$

where $\lambda^{(l)}_{\text{crit}} = \sum_{j\in\mathcal{J}}\gamma_j B^{(l)}_j$ and $(w)_+ = \max(0,w)$. We note that the above lower bound is tight only for $\lambda_l > \lambda^{(l)}_{\text{crit}}$, which represents the *heavy load* scenario. In Section 3.5, we will compare the blocking probability of jobs under the MV(d) scheme with the lower bound derived above to show the efficacy of the MV(d) scheme.

Rate of decay of the stationary tail probabilities: If we specialize (3.2) to the case where only a single class of jobs ($L=1$) requiring one unit of resource ($A=1$) arrive at the system, then we obtain

$$P_{k,j} - P_{k+1,j} = \frac{\lambda}{\gamma_j k} \left[\left(\sum_{i=1}^{j} \gamma_i P_{k-1+C_i-C_j,i} + \sum_{i=j+1}^{M} \gamma_i P_{k+C_i-C_j,i} \right)^d - \left(\sum_{i=1}^{j-1} \gamma_i P_{k-1+C_i-C_j,i} + \sum_{i=j}^{M} \gamma_i P_{k+C_i-C_j,i} \right)^d \right], \quad (3.5)$$

where the arrival rate of the jobs is $N\lambda$. Using (3.5) we can explicitly characterize the rate of decay of the probabilities $P_{k,j},\ k \in \{0, 1, \ldots, C_j\}$, which denote stationary tail probabilities of the number of jobs in progress at a type j server in the limiting system. This is done in the following proposition whose proof is similar to the proof of Theorem 2 of [49] for the homogeneous loss model.

Proposition 3.3.1. *Let $\left\{\bar{P}_k, 0 \le k \le C_M\right\}$ be defined as follows: $\bar{P}_k = 1$ for $0 \le k \le k_0$ and*

$$\bar{P}_k = \frac{\lambda^{d^{k-k_0-1}}}{(\lfloor\lambda\rfloor + k - k_0)(\lfloor\lambda\rfloor + k - k_0 - 1)^d \ldots (\lfloor\lambda\rfloor + 1)^{d^{k-k_0-1}}}, \quad (3.6)$$

for $k_0 + 1 \le k \le C_M$, 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, with each job requiring one unit of resource, we have

$$\sum_{j=1}^{M} \gamma_j P_{k+C_j-C_M,j} \le \bar{P}_k \text{ for } 0 \le k \le C_M \quad (3.7)$$

In particular, the average blocking probability of jobs is upper bounded as

$$P^{avg}_{blocking} = \left(\sum_{j \in \mathcal{J}} \gamma_j P_{C_j,j} \right)^d \le \bar{P}^d_{C_M} \quad (3.8)$$

Proof. From (3.5) we obtain

$$\sum_{j\in\mathcal{J}}(k+1+C_j-C_M)_+\gamma_j(P_{k+1+C_j-C_M,j}-P_{k+2+C_j-C_M,j})$$
$$=\lambda\left[\left(\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}\right)^d-\left(\sum_{j=1}^{M}\gamma_j P_{k+1+C_j-C_M,j}\right)^d\right], \quad (3.9)$$

where $0 \le k \le C_M - 1$ and $(w)_+ = \max(0, w)$. From (3.9) the following can be shown to hold for $0 \le k \le C_M - 1$ using backward induction starting at $k = C_M - 1$.

$$\sum_{j=1}^{M}(k+1+C_j-C_M)_+\gamma_j P_{k+1+C_j-C_M,j} \le \lambda\left(\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}\right)^d. \quad (3.10)$$

From (3.10) it is clear that for $C_M - C_1 \le k \le C_M - 1$ we have

$$\left(\sum_{j=1}^{M}\gamma_j P_{k+1+C_j-C_M,j}\right) \le \frac{\lambda}{k+(C_1-C_M)+1}\left(\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}\right)^d, \quad (3.11)$$

Now, for $0 \le k \le k_0$, $\bar{P}_k = 1 \ge \left(\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}\right)$ holds trivially. Assume that $\left(\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}\right) \le \bar{P}_k$ holds for some $k \ge k_0$. Using induction, we will now show that the inequality must hold for $k+1$. Since $k_0 > C_M - C_1$, we have from (3.10) that

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

The last equality follows from the definition of $\left\{\bar{P}_k, 0 \le k \le C_M\right\}$. This completes the proof. □

The above proposition shows that for $d \ge 2$ the quantity $\sum_{j=1}^{M}\gamma_j P_{k+C_j-C_M,j}$, which denotes the probability that a randomly sampled server can hold at most $C_M - k$ additional jobs, decreases with the increase in k at a rate much faster than that for $d = 1$. This shows the efficacy of assigning each incoming job by comparing the states a small number (> 1) of servers in the system at the arrival instant of the job.

3.4 Mean field analysis

In this section we provide detailed proofs of the main results given in Section 3.3 using mean field analysis. We first introduce the notation and mathematical framework required 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 $\underline{e}$ we denote the L-dimensional vector of all ones, i.e., $\underline{e} = (1, 1, \ldots, 1)$. Furthermore, for each $j \in \mathcal{J}$ we denote the space of probability distributions on $\mathcal{S}_j$ by $\mathcal{V}_j$, i.e.,

$$\mathcal{V}_j = \left\{ (g_{\underline{n}})_{\underline{n} \in \mathcal{S}_j} : g_{\underline{n}} \geq 0 \text{ for all } \underline{n} \in \mathcal{S}_j, \sum_{\underline{n} \in \mathcal{S}_j} g_{\underline{n}} = 1 \right\}. \tag{3.12}$$

We note that the space $\mathcal{V}_j$ of probability measures on the finite space $\mathcal{S}_j$ is convex and compact under any norm. The set of empirical probability distributions on $\mathcal{S}_j$ when the system size is N is denoted by $\mathcal{V}_j^{(N)}$, i.e.,

$$\mathcal{V}_j^{(N)} = \left\{ (g_{\underline{n}})_{\underline{n} \in \mathcal{S}_j} \in \mathcal{V}_j : N\gamma_j g_{\underline{n}} \in \mathbb{Z}_+ \right\}. \tag{3.13}$$

We will mainly be interested in the spaces $\mathcal{V} = \prod_{j \in \mathcal{J}} \mathcal{V}_j$ and $\mathcal{V}^{(N)} = \prod_{j \in \mathcal{J}} \mathcal{V}_j^{(N)}$, which are the Cartesian products of the spaces $\mathcal{V}_j$ and $\mathcal{V}_j^{(N)}$, respectively, over $j \in \mathcal{J}$. A point in the space $\mathcal{V}$ (or $\mathcal{V}^{(N)}$) 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 $\mathcal{V}_j$ ($\mathcal{V}_j^{(N)}$). We observe that the space $\mathcal{V}$ is compact since for each $j \in \mathcal{J}$ the space $\mathcal{V}_j$ is compact.

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

Lemma 3.4.1. *Let $\mathbf{u} \in \mathcal{V}^{(N)}$ be any state of the process $\mathbf{x}^{(N)}(\cdot)$ 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 $(\underline{k}, i) \neq (\underline{n}, j)$. The generator $\boldsymbol{\mathcal{G}}_N$ of the Markov process $\mathbf{x}^{(N)}(\cdot)$ acting on functions $\varphi : \mathcal{V}^{(N)} \to \mathbb{R}$ is given by*

$$\boldsymbol{\mathcal{G}}_N\varphi(\mathbf{u}) = N\sum_{j\in\mathcal{J}}\sum_{\underline{n}\in\mathcal{S}_j}\sum_{l\in\mathcal{L}}\Bigg[\lambda_l\frac{F(\underline{n}-\underline{e}_l,j,\mathbf{u})}{E(\underline{n}-\underline{e}_l,j,j,\mathbf{u})}\gamma_j u_{\underline{n}-\underline{e}_l,j}\left(\varphi\left(\mathbf{u}-\frac{\mathbf{e}(\underline{n}-\underline{e}_l,j)}{N\gamma_j}+\frac{\mathbf{e}(\underline{n},j)}{N\gamma_j}\right)\right.$$

$$\left.-\varphi(\mathbf{u})\right)+\gamma_j u_{\underline{n},j}n_l\left(\varphi\left(\mathbf{u}+\frac{\mathbf{e}(\underline{n}-\underline{e}_l,j)}{N\gamma_j}-\frac{\mathbf{e}(\underline{n},j)}{N\gamma_j}\right)-\varphi(\mathbf{u})\right)\Bigg] I_{\underline{n}-\underline{e}_l\in\mathcal{S}_j}, \quad (3.14)$$

where I denotes the indicator function and for $\underline{n}\in\mathcal{S}_j$ and $i,j\in\mathcal{J}$ we have

$$E(\underline{n},i,j,\mathbf{u}) = \gamma_i \sum_{\substack{\underline{n}'\in\mathcal{S}_i:\\ \underline{n}'\cdot\underline{A}=\underline{n}\cdot\underline{A}+C_i-C_j}} u_{\underline{n}',i} \quad (3.15)$$

$$G(\underline{n},i,j,\mathbf{u}) = \gamma_i \sum_{\substack{\underline{n}'\in\mathcal{S}_i:\\ \underline{n}'\cdot\underline{A}>\underline{n}\cdot\underline{A}+C_i-C_j}} u_{\underline{n}',i} \quad (3.16)$$

$$GE(\underline{n},i,j,\mathbf{u}) = G(\underline{n},i,j,\mathbf{u}) + E(\underline{n},i,j,\mathbf{u}). \quad (3.17)$$

and

$$F(\underline{n},j,\mathbf{u}) = \left(\sum_{i=1}^{j} GE(\underline{n},i,j,\mathbf{u}) + \sum_{i=j+1}^{M} G(\underline{n},i,j,\mathbf{u})\right)^d$$

$$-\left(\sum_{i=1}^{j-1} GE(\underline{n},i,j,\mathbf{u}) + \sum_{i=j}^{M} G(\underline{n},i,j,\mathbf{u})\right)^d. \quad (3.18)$$

Proof. We first consider the transition of the system from the state $\mathbf{u}\in\mathcal{V}^{(N)}$ at t^- to the state $\mathbf{u}-\frac{\mathbf{e}(\underline{n}-\underline{e}_l,j)}{N\gamma_j}+\frac{\mathbf{e}(\underline{n},j)}{N\gamma_j}$ 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 servers, sampled at the arrival instant, 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 $N\times E(\underline{n}-\underline{e}_l,j,j,u)$ servers of type j in states $\underline{n}-\underline{e}_l$ ans $\underline{n}'$, respectively, the probability of the above mentioned event is given by $\binom{k}{1}\gamma_j u_{\underline{n}-\underline{e}_l,j}E^{k-1}(\underline{n}-\underline{e}_l,j,j,u)$. 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 one 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 $\sum_{k=1}^{d} \binom{k}{1} \frac{1}{k} \gamma_j u_{\underline{n}-\underline{e}_l,j} E^{k-1}(\underline{n} - \underline{e}_l, j, j, u) \left(\sum_{i=1}^{j-1} GE(\underline{n}, i, j, \mathbf{u}) + \sum_{i=j}^{M} G(\underline{n}, i, j, \mathbf{u}) \right)^{d-k}$ 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

$$r\left(\mathbf{u} \to \mathbf{u} - \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} + \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}\right) = N\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \mathbf{u})}{E(\underline{n} - \underline{e}_l, j, j, \mathbf{u})} \gamma_j u_{\underline{n}-\underline{e}_l,j}. \tag{3.19}$$

Next, we consider the transition from the state $\mathbf{u} \in \mathcal{V}^{(N)}$ 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

$$r\left(\mathbf{u} \to \mathbf{u} + \frac{\mathbf{e}(\underline{n} - \underline{e}_l, j)}{N\gamma_j} - \frac{\mathbf{e}(\underline{n}, j)}{N\gamma_j}\right) = N\gamma_j u_{\underline{n},j} n_l \tag{3.20}$$

The expression (3.14) now follows directly from the definition of $\boldsymbol{\mathcal{G}}_N$. □

Using the generator $\boldsymbol{\mathcal{G}}_N$, we now show that that as $N \to \infty$, the sequence of processes $\left(\mathbf{x}^{(N)}(\cdot)\right)_N$ converges to a deterministic process.

Theorem 3.4.1. *If $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \mathcal{V}$ as $N \to \infty$, then $\mathbf{x}^{(N)}(\cdot) \Rightarrow \mathbf{x}(\cdot)$ as $N \to \infty$, where the process $\mathbf{x}(\cdot)$, taking values in the space $\mathcal{V}$ is given by the unique solution of the following system of differential equations*

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{3.21}$$

$$\dot{\mathbf{x}}(t) = \mathbf{h}(\mathbf{x}(t)). \tag{3.22}$$

The mapping $\mathbf{h}$ *above is given by*

$$h_{\underline{n},j}(\mathbf{x}) = \sum_{l\in\mathcal{L}} \left[\lambda_l \frac{F(\underline{n}-\underline{e}_l, j, \mathbf{x})}{E(\underline{n}-\underline{e}_l, j, j, \mathbf{x})} x_{\underline{n}-\underline{e}_l,j} - n_l x_{\underline{n},j}\right] I_{\underline{n}-\underline{e}_l\in\mathcal{S}_j} - \left[\lambda_l \frac{F(\underline{n}, j, \mathbf{x})}{E(\underline{n}, j, j, \mathbf{x})} x_{\underline{n},j} - (n_l+1) x_{\underline{n}+\underline{e}_l,j}\right] I_{\underline{n}+\underline{e}_l\in\mathcal{S}_j}, \tag{3.23}$$

where I *denotes indicator function and* $E(\underline{n}, i, j, \mathbf{u})$*,* $F(\underline{n}, j, \mathbf{u})$ *are as defined in Lemma 3.4.1.*

Proof. The first step is to show that there exists a unique process $\mathbf{x}(\cdot)$ satisfying (3.21)-(3.22) and taking values in the space $\mathcal{V}$. To see this, we first note that for all $\mathbf{x}\in\mathcal{V}$ and $j\in\mathcal{J}$, we have $\sum_{\underline{n}\in\mathcal{S}_j} h_{\underline{n},j}(\mathbf{x}) = 0$. This ensures that if $\sum_{\underline{n}\in\mathcal{S}_j} x_{\underline{n},j}(0) = 1$ for all $j\in\mathcal{J}$, then $\sum_{\underline{n}\in\mathcal{S}_j} x_{\underline{n},j}(t) = 1$ for all $j\in\mathcal{J}$ and all $t\geq 0$. Furthermore, we note from (3.23) that if for some $j\in\mathcal{J}$, $\underline{n}\in\mathcal{S}_j$, $t\geq 0$ we have $x_{\underline{n},j}(t) = 0$, then $h_{\underline{n},j}(\mathbf{x}(t)) \geq 0$. Thus, any solution of (3.21)-(3.22) starting from $\mathbf{u}_0\in\mathcal{V}$ always remains in the space $\mathcal{V}$. To show that there exists a unique solution to (3.21)-(3.22), it is sufficient to show that the mapping $\mathbf{h}$ on $\mathcal{V}$ is Lipschitz continuous, i.e., there exists $K>0$ such that $\|\mathbf{h}(\mathbf{u})-\mathbf{h}(\mathbf{v})\| \leq K\|\mathbf{u}-\mathbf{v}\|$, where $\|\cdot\|$ denotes the L_1-norm defined as $\|\mathbf{u}-\mathbf{v}\| = \sum_{j\in\mathcal{J}}\sum_{\underline{n}\in\mathcal{S}_j} |u_{\underline{n},j} - v_{\underline{n},j}|$ for all $\mathbf{u},\mathbf{v}\in\mathcal{V}$. But this is clearly satisfied with $K = 2B + 2\lambda d + 8\lambda S(d-1)$, where $S = \sum_{j\in\mathcal{J}} |\mathcal{S}_j|$ and $B = \max_{\underline{n}\in\mathcal{S}_j}\left(\sum_{l\in\mathcal{L}} n_l\right)$. Hence, for each initial point $\mathbf{u}_0\in\mathcal{V}$ there exists a unique process $\mathbf{x}(\cdot,\mathbf{u}_0)$ satisfying (3.21)-(3.22) with $\mathbf{x}(0) = \mathbf{u}_0$.

The next step is to show that the partial derivatives $\frac{\partial \mathbf{x}(t,\mathbf{u})}{\partial u_{n,j}}$, $\frac{\partial^2 \mathbf{x}(t,\mathbf{u})}{\partial {u_{n,j}}^2}$, and $\frac{\partial^2 \mathbf{x}(t,\mathbf{u})}{\partial u_{n,j}\partial u_{n',j'}}$ exist and are bounded for $\mathbf{u}\in\mathcal{V}$. This can shown using the same line of arguments as in the proof of Lemma 2.6.2 of Chapter 2. We therefore do not repeat them here.

Finally, the statement of the theorem follows by noting that $\boldsymbol{\mathcal{G}}_{\boldsymbol{N}}\varphi(\mathbf{u}) \to \frac{d}{dt}\varphi(\mathbf{x}(t,\mathbf{u}))|_{t=0}$ as $N\to\infty$ uniformly in $\mathbf{u}$ for all $\varphi:\mathcal{V}\to\mathbb{R}$ such that φ has bounded partial derivatives of first and second orders with respect to each component of its argument. □

The process $\mathbf{x}(\cdot)$, defined in the theorem above, is referred to as the *mean field limit* of the system. To emphasize the dependence of $\mathbf{x}(\cdot)$ on its initial value, we denote the process $\mathbf{x}(\cdot)$ started at $\mathbf{x}(0) = \mathbf{u}_0$ by $\mathbf{x}(\cdot,\mathbf{u}_0)$.

We now characterize the properties of the equilibrium points $\boldsymbol{\pi} = (\pi_{\underline{n},j}, \underline{n} \in \mathcal{S}_j, j \in \mathcal{J})$ of the process $\mathbf{x}(\cdot)$ satisfying $\mathbf{x}(t, \boldsymbol{\pi}) = \boldsymbol{\pi}$ for all $t \geq 0$. Clearly, for any such equilibrium point $\boldsymbol{\pi}$ we must have $h_{\underline{n},j}(\boldsymbol{\pi}) = 0$ for all $\underline{n} \in \mathcal{S}_j$ and $j \in \mathcal{J}$. Hence, from (3.23) we obtain

$$\sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n} - \underline{e}_l, j, \boldsymbol{\pi})}{E(\underline{n} - \underline{e}_l, j, j, \boldsymbol{\pi})} \pi_{\underline{n}-\underline{e}_l,j} - n_l \pi_{\underline{n},j} \right] I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = \sum_{l \in \mathcal{L}} \left[\lambda_l \frac{F(\underline{n}, j, \boldsymbol{\pi})}{E(\underline{n}, j, j, \boldsymbol{\pi})} \pi_{\underline{n},j} - (n_l + 1) \pi_{\underline{n}+\underline{e}_l,j} \right] I_{\underline{n}+\underline{e}_l \in \mathcal{S}_j}, \quad (3.24)$$

for $\underline{n} \in \mathcal{S}_j$ and $j \in \mathcal{J}$. In the next theorem, we show that there exists an equilibrium point $\boldsymbol{\pi}$ of the mean field $\mathbf{x}(\cdot)$ in the space $\mathcal{V}$.

Theorem 3.4.2. *There exists a equilibrium point* $\boldsymbol{\pi}$ *of the system* (3.21)-(3.22) *in the space* $\mathcal{V}$.

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

$$\lambda^{(l)}_{\underline{n},j}(\mathbf{x}) = \lambda_l \frac{F(\underline{n}, j, \mathbf{x})}{E(\underline{n}, j, j, \mathbf{x})} > 0. \quad (3.25)$$

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

$$\sum_{l \in \mathcal{L}} \left[\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}(\mathbf{x}) y_{\underline{n}-\underline{e}_l,j}(\mathbf{x}) - n_l y_{\underline{n},j}(\mathbf{x}) \right] I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = \sum_{l \in \mathcal{L}} \left[\lambda^{(l)}_{\underline{n},j}(\mathbf{x}) y_{\underline{n},j}(\mathbf{x}) - (n_l + 1) y_{\underline{n}+\underline{e}_l,j}(\mathbf{x}) \right] I_{\underline{n}+\underline{e}_l \in \mathcal{S}_j}, \text{ for } j \in \mathcal{J} \text{ and } \underline{n} \in \mathcal{S}_j \quad (3.26)$$

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

$$\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}(\mathbf{x}) y_{\underline{n}-\underline{e}_l,j}(\mathbf{x}) I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = n_l y_{\underline{n},j}(\mathbf{x}) \text{ for all } \underline{n} \in \mathcal{S}_j, j \in \mathcal{J}. \quad (3.27)$$

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 $\mathbf{y}(\mathbf{x}) \in \mathcal{V}$ for all $\mathbf{x} \in \mathcal{V}$. Furthermore, the map $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$, as defined above, is continuous on the space $\mathcal{V}$. Since $\mathcal{V}$ 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 (3.26), it is clear that any fixed point $\boldsymbol{\pi}$ of the map $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$ satisfies (3.24) and hence is an equilibrium point of the mean field $\mathbf{x}(\cdot)$. This proves the existence of an equilibrium point $\boldsymbol{\pi}$ in $\mathcal{V}$ of the mean field $\mathbf{x}(\cdot)$. □

Remark 15. We note that for each $\mathbf{x} \in \mathcal{V}$, the mapping $\mathbf{x} \mapsto \mathbf{y}(\mathbf{x})$ satisfies (3.27). Hence, the fixed point $\boldsymbol{\pi}$ of the above map satisfies

$$\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}(\boldsymbol{\pi})\pi_{\underline{n}-\underline{e}_l,j} I_{\underline{n}-\underline{e}_l \in \mathcal{S}_j} = n_l \pi_{\underline{n},j} \text{ for all } \underline{n} \in \mathcal{S}_j, j \in \mathcal{J}. \tag{3.28}$$

We will later show that the above equations correspond to the detailed balance equations for the servers in the large N limit.

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

Theorem 3.4.3. *For the single class case ($L = 1$), the mean field $\mathbf{x}(\cdot)$ has a unique globally asymptotically stable equilibrium point $\boldsymbol{\pi} \in \mathcal{V}$.*

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. We now proceed to show the global asymptotic stability of any equilibrium of the mean field $\mathbf{x}(\cdot)$ 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, $\mathcal{S}_j = \{0, 1, \ldots, C_j\}$, $\mathcal{V}_j = \left\{(g_n)_{n\in\mathcal{S}_j} : g_n \geq 0, \sum_{n\in\mathcal{S}_j} g_n = 1\right\}$, $\mathcal{V} = \prod_{j\in\mathcal{J}} \mathcal{V}_j$. In this case, the mean field $\mathbf{x}(\cdot) = (x_{n,j}(\cdot), n \in \mathcal{S}_j, j \in \mathcal{J}) \in \mathcal{V}$ satisfies the following system of differential equations (from (3.21)-(3.22))

$$\mathbf{x}(0) = \mathbf{u}_0, \tag{3.29}$$

$$\dot{\mathbf{x}}(t) = \mathbf{h}(\mathbf{x}(t)), \tag{3.30}$$

with $\mathbf{h}$ is defined as

$$h_{n,j}(\mathbf{x}) = \left[\frac{\lambda}{\gamma_j} F(n-1,j,\mathbf{x}) - n x_{n,j}\right] I_{1\le n\le C_j} - \left[\frac{\lambda}{\gamma_j} F(n,j,\mathbf{x}) - (n+1) x_{n+1,j}\right] I_{0\le n\le C_j-1}. \tag{3.31}$$

The mean field can be equivalently 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 (3.30) and (3.31) we have

$$\tilde{\mathbf{x}}(0) = \tilde{\mathbf{u}}_0, \tag{3.32}$$

$$\dot{\tilde{\mathbf{x}}}(t) = \tilde{\mathbf{h}}(\tilde{\mathbf{x}}(t)), \tag{3.33}$$

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

$$\tilde{h}_{k,j}(\tilde{\mathbf{x}}) = \frac{\lambda}{\gamma_j}\left[\left(\sum_{i=1}^{j} \gamma_i \tilde{x}_{k-1+C_i-C_j,i} + \sum_{i=j+1}^{M} \gamma_i \tilde{x}_{k+C_i-C_j,i}\right)^d - \left(\sum_{i=1}^{j-1} \gamma_i \tilde{x}_{k-1+C_i-C_j,i} + \sum_{i=j}^{M} \gamma_i \tilde{x}_{k+C_i-C_j,i}\right)^d\right] - k\left(\tilde{x}_{k,j} - \tilde{x}_{k+1,j}\right) \tag{3.34}$$

We say that $\tilde{\mathbf{u}} \le \tilde{\mathbf{u}}'$ if $\tilde{u}_{k,j} \le \tilde{u}'_{k,j}$ 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 3.4.2. *If $\tilde{\mathbf{u}}_0 \le \tilde{\mathbf{u}}'_0$ then $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0) \le \tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}'_0)$ for all $t \ge 0$.*

Proof. Clearly, the right hand side of (3.34) is non-decreasing in $\tilde{x}_{n,i}(t)$ for all $(n,i) \ne (k,j)$. Hence, (3.34) defines a quasi-monotone system of differential equations. The proof of the lemma now follows directly from pages 70-74 of [38]. □

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 $\tilde{\mathbf{u}}_0$. From (3.34) we obtain

$$\frac{dz(t,\tilde{\mathbf{u}}_0)}{dt} = \lambda\left(1-\left(\sum_{j\in\mathcal{J}}\gamma_j\tilde{x}_{C_j,j}(t,\tilde{\mathbf{u}}_0)\right)^d\right) - z(t,\tilde{\mathbf{u}}_0). \tag{3.35}$$

Let $\tilde{\boldsymbol{\pi}}$ be an equilibrium point of the process $\tilde{\mathbf{x}}(\cdot)$. From (3.35) we have

$$\lambda\left(1-\left(\sum_{j\in\mathcal{J}}\gamma_j\tilde{\pi}_{C_j,j}\right)^d\right) = z(t,\boldsymbol{\pi}) = \sum_{j\mathcal{J}}\gamma_j\sum_{k=1}^{C_j}\tilde{\pi}_{k,j} \tag{3.36}$$

Now, from Lemma 3.4.2 we have

$$\tilde{\mathbf{x}}(t,\min(\tilde{\mathbf{u}}_0,\tilde{\boldsymbol{\pi}})) \le \tilde{\mathbf{x}}(t,\tilde{\mathbf{u}}_0) \le \tilde{\mathbf{x}}(t,\max(\tilde{\mathbf{u}}_0,\tilde{\boldsymbol{\pi}})), \tag{3.37}$$

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

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

$$\int_0^\infty (\tilde{x}_{n,j}(t,\tilde{\mathbf{u}}_0) - \tilde{\pi}_{n,j})\,dt < \infty,\ \text{for all } j\in\mathcal{J}, 1\le n\le C_j. \tag{3.38}$$

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

$$\int_0^\infty (\tilde{\pi}_{n,j} - \tilde{x}_{n,j}(t,\tilde{\mathbf{u}}_0))\,dt < \infty,\ \text{for all } j\in\mathcal{J}, 1\le n\le C_j \tag{3.39}$$

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

For $\tilde{\mathbf{u}}_0 \ge \tilde{\boldsymbol{\pi}}$ we have from Lemma 3.4.2 that $\tilde{\mathbf{x}}(t,\tilde{\mathbf{u}}_0) \ge \tilde{\boldsymbol{\pi}}$ for all $t \ge 0$. Hence, to prove (3.38) 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

$$
\int_0^\tau \left(z(t, \tilde{\mathbf{u}}_0) - \sum_{j\in\mathcal{J}} \gamma_j \sum_{n=1}^{C_j} \tilde{\pi}_{n,j} \right) dt = -\int_0^\tau \frac{dz(t,\tilde{\mathbf{u}}_0)}{dt} dt
$$

$$
-\int_0^\tau \lambda \left(\left(\sum_{j\in\mathcal{J}} \gamma_j \tilde{\pi}_{C_j,j} \right)^d - \left(\sum_{j\in\mathcal{J}} \gamma_j \tilde{x}_{C_j,j}(t, \tilde{\mathbf{u}}_0) \right)^d \right) dt
$$

$$
\leq (z(0, \tilde{\mathbf{u}}_0) - z(\tau, \tilde{\mathbf{u}}_0)) \leq z(0, \tilde{\mathbf{u}}_0)
$$

where the first equality follows from (3.35) and (3.36); 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 $\tau \to \infty$. This completes the proof. □

We now consider uniqueness and stability of the equilibrium point for the multi-class case, where the monotonicity property, similar to the one established in Lemma 3.4.2, does not hold [56]. In this case, it is difficult to show the global asymptotic stability of the equilibrium point due to the lack of the monotonicity property. However, in the following theorem we show that a unique equilibrium point $\boldsymbol{\pi} \in \mathcal{V}$ of the mean field $\mathbf{x}(\cdot)$ exists and it is asymptotically stable, i.e., there exists $\epsilon > 0$ such that for any $\mathbf{u}_0 \in \mathcal{V}$ satisfying $\|\mathbf{u}_0 - \boldsymbol{\pi}\|_2 < \epsilon$, we have $\mathbf{x}(t, \mathbf{u}_0) \to \boldsymbol{\pi}$ as $t \to \infty$. This implies that if the initial condition for the process $\mathbf{x}^{(N)}(\cdot)$ is chosen sufficiently close to the equilibrium point, then the corresponding mean field limit converges asymptotically to its unique equilibrium point.

Theorem 3.4.4. *For the multi-class case ($L > 1$) there exists a unique equilibrium point $\boldsymbol{\pi} \in \mathcal{V}$ of the mean field $\mathbf{x}(\cdot)$. Furthermore, the equilibrium point $\boldsymbol{\pi}$ is asymptotically stable.*

Proof sketch: To prove the theorem, we first express the mean field in terms of the tail sums $\tilde{x}_{\underline{k},j} = \sum_{\underline{n}\in\mathcal{S}_j : \underline{n}.\underline{A} \geq \underline{k}.\underline{A}} x_{\underline{n},j}$, $k \in \mathcal{S}_j$ assuming without loss of generality that the vector $\underline{A}$ is such that for any two states $\underline{n}, \underline{n}' \in \mathcal{S}_j$ we have $\underline{n}.\underline{A} \neq \underline{n}'.\underline{A}$ [1]. We note that $\tilde{x}_{\underline{0},j}(t) = 1$ for all $j \in \mathcal{J}$, $t \geq 0$. Therefore, expressed in terms of the tail sums, the mean field is given by

[1]In case this does not hold, we can order the states of the servers of each type in the increasing order of their resource requirements. States having the same resource requirement can be ordered arbitrarily. The tail sums can then be defined according to the ordering of the states.

$$\tilde{\mathbf{x}}(0) = \tilde{\mathbf{u}}_0, \tag{3.40}$$

$$\dot{\tilde{x}}_{\underline{k},j}(t) = \tilde{h}_{\underline{k},j}(\tilde{\mathbf{x}}(t)), \tag{3.41}$$

where $\tilde{\mathbf{x}} = (\tilde{x}_{\underline{k},j}, \underline{k} \in \mathcal{S}_j \backslash \underline{0}, j \in \mathcal{J})$ and $\tilde{h}_{\underline{k},j} = \sum_{\underline{n} \in \mathcal{S}_j : \underline{n}.\underline{A} \geq \underline{k}.\underline{A}} h_{\underline{n},j}$ for $\underline{k} \in \mathcal{S}_j \backslash \underline{0}$, $j \in \mathcal{J}$. We define the mapping $\tilde{\mathbf{h}} = (\tilde{h}_{\underline{k},j}, \underline{k} \in \mathcal{S}_j \backslash \underline{0}, j \in \mathcal{J})$. The next (second step) is to verify that the mapping $\tilde{\mathbf{h}}$, when seen as a mapping from $\mathbb{R}^{\sum_{j \in \mathcal{J}} |\mathcal{S}_j| - M}$ to itself, is proper, i.e., $\|\tilde{\mathbf{h}}(\tilde{\mathbf{x}})\|_2 \to \infty$ if $\|\tilde{\mathbf{x}}\|_2 \to \infty$. Finally, the third step is to verify that the Jacobian matrix $\tilde{\mathbf{J}}(\tilde{\mathbf{x}})$ of $\tilde{\mathbf{h}}$ evaluated at $\tilde{\mathbf{x}}$ is negative definite (and hence non-singular) for all $\tilde{\mathbf{x}} \in \mathbb{R}^{\sum_{j \in \mathcal{J}} |\mathcal{S}_j| - M}$. The third step shows that the mapping $\tilde{\mathbf{h}}$ locally homeomorphic at every point in $\mathbb{R}^{\sum_{j \in \mathcal{J}} |\mathcal{S}_j| - M}$ [[57], Theorem 3.1.5, Page 113]. According to the Hadamard's global inverse function theorem [[57], Theorem 5.1.4 (i), Page 221], the second and the third step will then together imply that the mapping $\tilde{\mathbf{h}}$ is globally homeomorphic on $\mathbb{R}^{\sum_{j \in \mathcal{J}} |\mathcal{S}_j| - M}$, i.e, the inverse exists and is continuous at every point on $\mathbb{R}^{\sum_{j \in \mathcal{J}} |\mathcal{S}_j| - M}$. This implies in particular that $\tilde{\mathbf{h}}^{-1}(\mathbf{0})$ is unique proving the uniqueness of the equilibrium point of the mean field. The negative definiteness of the Jacobian matrix, as established in the third step, will imply that the mean field $\tilde{\mathbf{x}}(t, \tilde{\mathbf{u}}_0)$ converges asymptotically to its unique equilibrium point $\tilde{\boldsymbol{\pi}}$ for all $\mathbf{u}_0$ sufficiently close to the equilibrium point $\boldsymbol{\pi}$.

The steps described above can be shown for any $L > 1$ (multiple classes) and $M > 1$ (heterogeneous servers). However, for notational convenience, we show the steps for a particular example which is representative of the general case. The proof for the general case is identical to the proof discussed below with necessary alterations in the notations.

We choose a system with parameters: $L = 2$, $M = 2$, $d = 2$ $C_1 = 3$, $C_2 = 4$, $A_1 = 2$, $A_2 = 3$. For the above parameter setting we have $\mathcal{S}_1 = \{(0,0), (1,0), (0,1)\}$ and $\mathcal{S}_2 = \{(0,0), (1,0), (0,1), (2,0)\}$. In this case the mean field can be expressed in terms of the vector of tail sums $\tilde{\mathbf{x}} = (\tilde{x}_{(1,0),1}, \tilde{x}_{(0,1),1}, \tilde{x}_{(1,0),2}, \tilde{x}_{(0,1),2}, \tilde{x}_{(2,0),2})$ as follows (we omit t for convenience)

$$
\begin{aligned}
\tilde{x}_{(0,0),1} &= 1, && (3.42)\\
\frac{d\tilde{x}_{(1,0),1}}{dt} &= h_{(1,0),1}(\tilde{\mathbf{x}}) = \frac{\lambda_1+\lambda_2}{\gamma_1}\vartheta_1(\tilde{\mathbf{x}}) - \tilde{x}_{(1,0),1} && (3.43)\\
\frac{d\tilde{x}_{(0,1),1}}{dt} &= h_{(0,1),1}(\tilde{\mathbf{x}}) = \frac{\lambda_2}{\gamma_1}\vartheta_1(\tilde{\mathbf{x}}) - \tilde{x}_{(0,1),1} && (3.44)\\
\tilde{x}_{(0,0),2} &= 1, && (3.45)\\
\frac{d\tilde{x}_{(1,0),2}}{dt} &= h_{(1,0),2}(\tilde{\mathbf{x}}) = \frac{\lambda_1+\lambda_2}{\gamma_2}\vartheta_2(\tilde{\mathbf{x}}) - \tilde{x}_{(1,0),2} + \tilde{x}_{(2,0),2} && (3.46)\\
\frac{d\tilde{x}_{(0,1),2}}{dt} &= h_{(0,1),2}(\tilde{\mathbf{x}}) = \frac{\lambda_2}{\gamma_2}\vartheta_2(\tilde{\mathbf{x}}) + \frac{\lambda_1}{\gamma_2}\vartheta_3(\tilde{\mathbf{x}}) - \tilde{x}_{(0,1),2} - \tilde{x}_{(2,0),2} && (3.47)\\
\frac{d\tilde{x}_{(2,0),2}}{dt} &= h_{(2,0),2}(\tilde{\mathbf{x}}) = \frac{\lambda_1}{\gamma_2}\vartheta_3(\tilde{\mathbf{x}}) - 2\tilde{x}_{(2,0),2}, && (3.48)
\end{aligned}
$$

where

$$
\begin{aligned}
\vartheta_1(\tilde{\mathbf{x}}) &= (\gamma_1 + \gamma_2\tilde{x}_{(1,0),2})^2 - (\gamma_1\tilde{x}_{(1,0),1} + \gamma_2\tilde{x}_{(1,0),2})^2, && (3.49)\\
\vartheta_2(\tilde{\mathbf{x}}) &= 1 - (\gamma_1 + \gamma_2\tilde{x}_{(1,0),2})^2, && (3.50)\\
\vartheta_3(\tilde{\mathbf{x}}) &= (\gamma_1\tilde{x}_{(1,0),1} + \gamma_2\tilde{x}_{(1,0),2})^2 - (\gamma_1\tilde{x}_{(1,0),1} + \gamma_2\tilde{x}_{(0,1),2})^2. && (3.51)
\end{aligned}
$$

Hence, each component of the mapping $\tilde{\mathbf{h}} = (h_{(1,0),1}, h_{(0,1),1}, h_{(1,0),2}, h_{(0,1),2}, h_{(2,0),2})$ is a polynomial on $\mathbb{R}^5$. It is easy to see from the expressions of the polynomials that if any subset of components of the vector $\tilde{\mathbf{x}}$ approaches to ∞, then at least one of the components of $\tilde{\mathbf{h}}$ approaches to ∞ or $-\infty$. Therefore, $\tilde{\mathbf{h}}$ is proper on $\mathbb{R}^5$ to itself. Finally, the Jacobian matrix $\tilde{\mathbf{J}}(\tilde{x})$ of the map $\tilde{\mathbf{h}}$ computed at any point $\tilde{x} \in \mathbb{R}^5$ is given by

$$
\tilde{\mathbf{J}}(\tilde{\mathbf{x}}) = \begin{bmatrix}
-2(\lambda_1+\lambda_2)(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(1,0),2})-1 & 0 & 2(\lambda_1+\lambda_2)\gamma_2(1-\tilde{x}_{(1,0),1}) & 0 & 0\\
-2\lambda_2(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(1,0),2}) & -1 & 2\lambda_2\gamma_2(1-\tilde{x}_{(1,0),1}) & 0 & 0\\
0 & 0 & -2(\lambda_1+\lambda_2)(\gamma_1+\gamma_2\tilde{x}_{(1,0),2})-1 & 0 & 1\\
2\lambda_1\gamma_1(\tilde{x}_{(1,0),2}-\tilde{x}_{(0,1),2}) & 0 & -2\lambda_2(\gamma_1+\gamma_2\tilde{x}_{(1,0),2})+2\lambda_1(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(1,0),2}) & -2\lambda_1(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(1,0),2})-1 & -1\\
2\lambda_1\gamma_1(\tilde{x}_{(1,0),2}-\tilde{x}_{(0,1),2}) & 0 & 2\lambda_1(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(1,0),2}) & -2\lambda_1(\gamma_1\tilde{x}_{(1,0),1}+\gamma_2\tilde{x}_{(0,1),2}) & -2
\end{bmatrix}
$$

A Routh-Hurwitz test [58] of the characteristic polynomial of the matrix $\tilde{\mathbf{J}}(\tilde{\mathbf{x}})$ shows that all the eigenvalues of the matrix have strictly negative real parts. Therefore, $\tilde{\mathbf{J}}(\tilde{\mathbf{x}})$ is negative

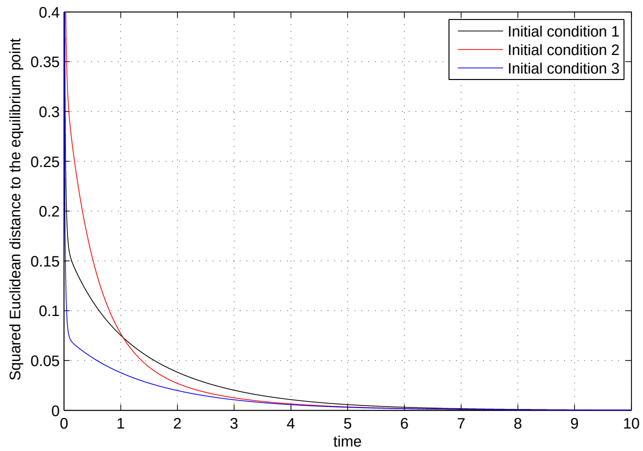

Figure 3.1: Squared Euclidean distance between the mean field and its equilibrium point as a function of time for random initial states. Parameters: $M = 2$, $L = 2$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 3$, $C_2 = 4$, $A_1 = 2$, $A_2 = 3$, $\lambda_1 = 20$, $\lambda_2 = 40$.

definite (hence non-singular) everywhere in $\mathbb{R}^5$. Thus, we conclude that the system has a unique equilibrium point $\tilde{\boldsymbol{\pi}}$ and it is asymptotically stable. $\square$

Remark 16. Although we have only established local asymptotic stability of the equilibrium point $\boldsymbol{\pi}$ for the multi-class case, our numerical results suggest that $\boldsymbol{\pi}$ is globally asymptotically stable, i.e., the convergence $\tilde{\mathbf{x}}(t) \to \tilde{\boldsymbol{\pi}}$ holds for all $\tilde{\mathbf{x}}(0)$ in the space of tail sums. As an example, in Figure 3.1 we plot the squared Euclidean distance between the vector $\tilde{\mathbf{x}} = (\tilde{x}_{(1,0),1}, \tilde{x}_{(0,1),1}, \tilde{x}_{(1,0),2}, \tilde{x}_{(0,1),2}, \tilde{x}_{(2,0),2})$ and the equilibrium point $\tilde{\boldsymbol{\pi}} = (\tilde{\pi}_{(1,0),1}, \tilde{\pi}_{(0,1),1}, \tilde{\pi}_{(1,0),2}, \tilde{\pi}_{(0,1),2}, \tilde{\pi}_{(2,0),2})$ as a function of time for different initial conditions. The parameters are chosen as $M = 2, L = 2, C_1 = 3, C_2 = 4, A_1 = 2, A - 2 = 3, \gamma_1 = \gamma_2 = 0.5, \lambda_1 = 20, \lambda_2 = 40$. We observe that for all initial conditions the squared Euclidean distance converges to 0 as $t \to \infty$.

The global asymptotic stability for the mean field corresponding to the multi-class case is also suggested by the Markus-Yamabe conjecture [59] which states that if a continuously differentiable map $\tilde{\mathbf{h}}$ on $\mathbb{R}^n$ has a unique equilibrium point and its Jacobian matrix is negative definite at each point in $\mathbb{R}^n$ then the equilibrium point is globally asymptotically stable. Although this conjecture has been proven only for $n = 2$ and is in general not true for $n > 2$, the exceptions are few and pathological.

We note that for each N the process $\mathbf{x}^{(N)}(\cdot)$ is positive recurrent and hence has a unique stationary distribution. Let $\mathbf{x}^{(N)}(\infty)$ denote the steady state of the process $\mathbf{x}^N(\cdot)$. In the next theorem, we show that $\mathbf{x}^{(N)}(\infty)$ concentrates near $\boldsymbol{\pi}$ as $N \to \infty$.

Theorem 3.4.5. *We have* $\mathbf{x}^{(N)}(\infty) \Rightarrow \boldsymbol{\pi}$ *as* $N \to \infty$.

Proof. The proof is essentially the same as that of Theorem 2.6.5 of Chapter 2. We first note that the space $\mathcal{V}$ is compact. Hence, the sequence of probability measures $(\text{Law}\left(\mathbf{x}^{(N)}(\infty)\right))_N$ on $\mathcal{V}$ is tight. According to Theorem C.1.2 this implies that the sequence $\left(\text{Law}\left(\mathbf{x}^{(N)}(\infty)\right)\right)_N$ has limit points. We now show that all the limit points coincide with $\delta_{\boldsymbol{\pi}}$. Due to Theorem 3.4.1, any limit point of the sequence $\left(\text{Law}\left(\mathbf{x}^{(N)}(\infty)\right)\right)_N$ must be invariant distribution of the map $\mathbf{u}_0 \mapsto \mathbf{x}(t, \mathbf{u}_0)$. Since by Theorem 3.4.2 there exists a unique invariant distribution $\delta_{\boldsymbol{\pi}}$ of the map $\mathbf{u}_0 \mapsto \mathbf{x}(t, \mathbf{u}_0)$ in $\mathcal{V}$ (see Remark 11), we conclude that every limit point of $\left(\text{Law}\left(\mathbf{x}^{(N)}(\infty)\right)\right)_N$ must coincide with $\delta_{\boldsymbol{\pi}}$, or equivalently $\mathbf{x}^{(N)}(\infty) \Rightarrow \boldsymbol{\pi}$ as $N \to \infty$. This completes the proof. □

3.4.1 Propagation of chaos

So far we have considered the convergence of the process $\mathbf{x}^{(N)}(\cdot)$ that describes the evolution of the empirical distribution of the server states. We now focus on servers in a given finite set of tagged servers. We show that as the system size grows the tagged servers become independent of each other, provided that the initial state of the overall system satisfies the intra-type exchangeability criterion as defined in 2.6.2. As discussed earlier, this property is formally known as the *asymptotic independence property* or the *propagation of chaos* property. We show that the stationary distribution of state of each server in the limiting system is determined by the unique stationary point $\boldsymbol{\pi}$ of the the process $\mathbf{x}(\cdot)$. To formally state our 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 $k \in \{1, 2, \ldots, N\gamma_j\}$, $j \in \mathcal{J}$.
- For each $j \in \mathcal{J}$ and $t \geq 0$, we denote by $x_j(t)$, the distribution on $\mathcal{S}_j$ given by $x_j(t) = (x_{\underline{n},j}(t), n \in \mathcal{S}_j)$, where $\mathbf{x}(\cdot)$ is the mean field process satisfying (3.21)-(3.22). Furthermore, we define $x_j(\infty) = (\pi_{\underline{n},j}, n \in \mathcal{S}_j)$.

Theorem 3.4.6. *For the model considered in Section 3.2, let* $(q_{k,j}^{(N)}(0), 1 \leq k \leq N\gamma_j, 1 \leq j \leq M)$ *be intra-type exchangeable in the sense of Definition 2.6.2 and let* $\mathbf{x}^{(N)}(0) \Rightarrow \mathbf{u}_0 \in \mathcal{V}$ *as* $N \to \infty$. *Then the following holds*

1. *For each $j \in \mathcal{J}$, $1 \leq k \leq N\gamma_j$, and $t \in [0,\infty]$, $q_{k,j}^{(N)}(t) \Rightarrow U_j(t)$ as $N \to \infty$, where $U_j(t)$ is a random variable with distribution $x_j(t)$.*

2. *Fix positive integers $1 \leq r_1 \leq N\gamma_1, 1 \leq r_2 \leq N\gamma_2, \ldots, 1 \leq r_M \leq N\gamma_M$. Then, for each $t \in [0,\infty]$,*

$$\left(q_{k,j}^{(N)}(t), 1 \leq k \leq r_j, 1 \leq j \leq M\right) \Rightarrow \left(U_{k,j}(t), 1 \leq k \leq r_j, 1 \leq j \leq M\right), \tag{3.52}$$

as $N \to \infty$, 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 $x_j(t)$ for all $1 \leq k \leq r_j$.

Proof. With the observation that, under the MV(d) scheme, servers of the same type are statistically indistinguishable the proof is identical to that of Theorem 2.6.6. □

The above theorem shows that the servers in the limiting system are independent of each other and the stationary probability that a server of type $j \in \mathcal{J}$ is in state $\underline{n} \in \mathcal{S}_j$ is given by $\pi_{\underline{n},j}$. 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 $\mathbf{x} \in \mathcal{V}$ is as defined in (3.25).

Proposition 3.4.1. *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*

$$\lambda_{\underline{n},j}^{(l)}(\boldsymbol{\pi}) = \lambda_l \frac{F(\underline{n}, j, \boldsymbol{\pi})}{E(\underline{n}, j, j, \boldsymbol{\pi})}, \tag{3.53}$$

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 3.4.1.

Proof. We consider a *tagged* server of type j and the class-l arrivals that have the tagged server as one of their 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 $\frac{d}{N} \times N\lambda_l = d\lambda_l$.

Now 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 at the tagged server actually joins the tagged server depending on the states of the other possible destination servers. However, as $N \to \infty$, due to the asymptotic independence property established in Theorem 3.4.6, 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 $N \to \infty$. Now the arrival rates of this Poisson process can be obtained using arguments similar to those in the proof of Proposition 2.6.3. We see that the arrival rates are given by (3.53). □

From Remark 15, we already know that the equilibrium point $\boldsymbol{\pi}$ satisfies

$$\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}(\boldsymbol{\pi})\pi_{\underline{n}-\underline{e}_l,j}I_{\underline{n}-\underline{e}_l\in\mathcal{S}_j} = n_l\pi_{\underline{n},j} \text{ for } \underline{n} \in \mathcal{S}_j \text{ and } l \in \mathcal{L}. \tag{3.54}$$

Now, since (by Proposition 3.4.1) $\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}(\boldsymbol{\pi})$ 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 $\underline{n} \in \mathcal{S}_j$ 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 [60, 61] at each server in the limiting system. This allows efficient computation of the blocking probabilities for each class of jobs.

Proposition 3.4.2. *Let $P_{k,j}$, for $1 \le k \le C_j$ and $j \in \mathcal{J}$, denote the stationary probability that a server in the limiting system has at least k units of occupied resource, i.e., $P_{k,j} = \sum_{\substack{\underline{n}\in\mathcal{S}_j:\\ \underline{n}.\underline{A}\ge k}} \pi_{\underline{n},j}$. Then $P_{k,j}$ satisfies (3.2) for $0 \le k \le C_j - 1$, where $P_{k,j} = 1$ for $k \le 0$, and $P_{C_j+1,j} = 0$ for all $j \in \mathcal{J}$.*

Proof. For $j \in \mathcal{J}$, $0 \le k \le C_j$, we define the set $\mathcal{D}_{k,j}$ as follows:

$$\mathcal{D}_{k,j} = \{\underline{n} \in \mathcal{S}_j : \underline{n} \cdot \underline{A} = k\}. \tag{3.55}$$

Thus, $\mathcal{D}_{k,j}$ denotes the set of states in $\mathcal{S}_j$ for which the total number of occupied resource units 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 using (3.53) that

$$\lambda^{(l)}_{\underline{n}-\underline{e}_l,j} = \frac{\lambda_l}{\gamma_j\left(P_{k-A_l,j} - P_{k-A_l,j}\right)}\left[\left(\sum_{i=1}^{j}\gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j+1}^{M}\gamma_i P_{k-A_l+C_i-C_j+1,i}\right)^d - \left(\sum_{i=1}^{j-1}\gamma_i P_{k-A_l+C_i-C_j,i} + \sum_{i=j}^{M}\gamma_i P_{k-A_l+C_i-C_j+1,i}\right)^d\right]. \quad (3.56)$$

Now multiplying (3.54) by A_l and summing over all $\underline{n} \in \mathcal{D}_{k,j}$ and all $l \in \mathcal{L}$ we have

$$\sum_{l\in\mathcal{L}}\sum_{\underline{n}\in\mathcal{D}_{k,j}} A_l\lambda^{(l)}_{\underline{n}-\underline{e}_l,j}\pi_{\underline{n}-\underline{e}_l,j}I_{\underline{n}-\underline{e}_l\in\mathcal{S}_j} = \sum_{l\in\mathcal{L}}\sum_{\underline{n}\in\mathcal{D}_{k,j}} n_l A_l \pi_{\underline{n},j} \text{ for } \underline{n}\in\mathcal{S}^j \text{ and } l\in\mathcal{L} \quad (3.57)$$

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

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

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

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

This completes the proof. □

Remark 17. 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 [34]. 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 3.4.1 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 (3.53). This implies that the detailed balance equations given by (3.54) also hold for general service time distributions. Since the servers in the system are loss servers, the detailed balance condition implies that that the stationary distribution of each server in the limiting system is *insensitive* to service time distributions (see Theorem 1 of [40]). We refer to this property as the *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.

3.5 Numerical Results

In this section, we first present simulation results that indicate the accuracy of the mean field analysis of the MV(d) scheme in predicting the performance of the scheme for large but finite systems. To show the efficacy of the MV(d) scheme, we then compare the average blocking probability of jobs under the MV(d) scheme with that under the randomized static schemes. We also compare the average blocking probability of jobs under the MV(d) scheme with the lower bound given in (3.4) to show that the MV(d) scheme is nearly optimal in terms of reducing the average blocking probability of jobs. Finally, we provide results that support insensitivity of the MV(d) scheme in the large system limit. All simulation results presented in this section are the averages of 10,000 independent runs.

To investigate the accuracy of the asymptotic analysis of the MV(d) scheme, we compare the average blocking probability obtained from the mean field analysis to that obtained by simulating the system for finite values of N. We set the following parameter values: $L = 1$, $A_1 = 1$, $M = 2$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 20$, $C_2 = 25$, $\lambda_1 = 30$, i.e., we consider the single class case where every job requires one unit of resource from each server. In this case, the blocking probability can be obtained using (3.5). The results are shown in Table 3.1. We observe that the difference between the asymptotic results and the

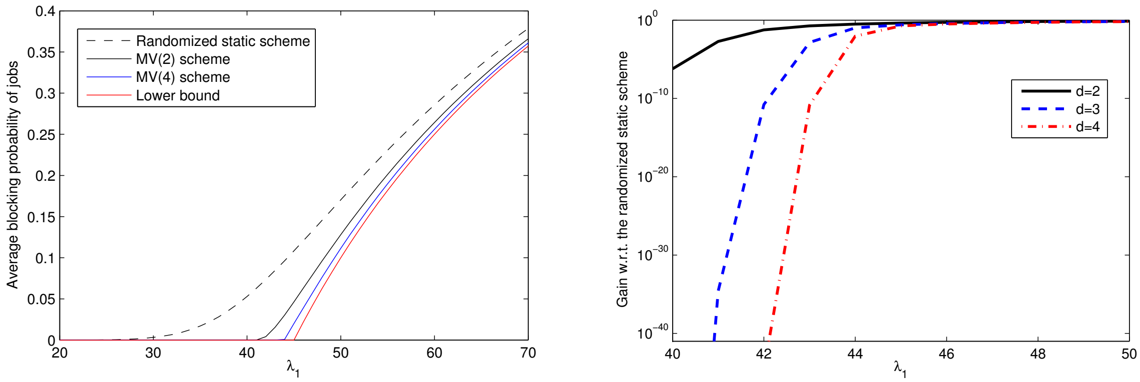

(a) Average blocking probability as a function of λ_1. (b) Ratio of average blocking probability under the MV(d) scheme to that under the randomized static scheme as a function of λ_1

Figure 3.2: Efficacy of the MV(d) scheme

corresponding simulation results decreases with the increase in N. Furthermore, for the same value of N, increasing d increases percentage of deviation between the analytical and simulation results. This is expected since for finite N increasing d increases the correlation between the server states which acts in opposition to the independence of servers in the limiting system. From our observations we conclude that the mean field analysis quite accurately captures the behavior of the finite system operating under the MV(d) scheme.

Table 3.1: Accuracy of the mean field analysis of the MV(d) scheme

d	Asymptotic	$N = 20$	$N = 50$	$N = 100$	$N = 200$
2	0.2751	0.2761	0.2757	0.2753	0.2751
4	0.2616	0.2651	0.2618	0.2616	0.2616
6	0.2576	0.2593	0.2588	0.2580	0.2577
8	0.2557	0.2584	0.2578	0.2569	0.2562
10	0.2545	0.2574	0.2565	0.2559	0.2549

In Figure 3.2, we compare average blocking probability of jobs under the MV(d) scheme with that under the randomized static scheme in which an incoming job is assigned to a server of type j with probability p_j independent of the states of the servers. We set $p_j =$

$\gamma_j C_j / \sum_{i=1}^{M} \gamma_i C_i$ since in this case p_j is proportional to both γ_j and C_j.[2] The parameters are chosen to be $L = 1$, $M = 2$, $N = 100$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 30$, and $C_2 = 60$. For this parameter setting, the critical load (see (3.4)) is given by $\lambda_{\text{crit}}^{(1)} = \gamma_1 C_1 + \gamma_2 C_2 = 45$. In Figure 3.2(a), we plot the average blocking probability of jobs as a function of λ_1 for the MV(d) scheme and the randomized static scheme. We have also plotted the lower bound obtained from (3.4). We observe that the average blocking probability of jobs under the MV(d) scheme is significantly lower than that under the randomized static scheme. We also observe that with the increase in d the average blocking probability reduces. Furthermore, the average blocking probability for $d = 4$ is seen to be quite close the lower bound for $\lambda_1 > \lambda_{\text{crit}}^{(1)}$. Hence, we conclude that the the MV(d) scheme is nearly optimal in terms of reducing the average blocking for probability even for small d.

In Figure 3.2(b), we plot the ratio of the average blocking probability under the MV(d) scheme to that under the randomized static scheme as a function of λ_1. The y-axis is drawn in the log scale. It is clear that the average blocking probability under the MV(d) scheme is orders-of-magnitude lower than that under the randomized static scheme. This shows the advantage of comparing the states of a small number of randomly sampled servers to assign each incoming request over randomly assigning the job requests to the servers.

So far we have investigated single class systems. We now numerically study a multi-class system with the following parameters: $N = 200, L = 2, \lambda_1 = 0.2\lambda, \lambda_2 = 0.8\lambda, C_1 = 10, C_2 = 20, \gamma_1 = \gamma_2 = 0.5, A_1 = 5, A_2 = 2$, where $\lambda = \lambda_1 + \lambda_2$ denotes the total normalized (by the number of servers) arrival rate of jobs, all classed combined. In Table 3.2, we compare the average blocking probability of each class of jobs under the MV(2) scheme with that under the randomized static scheme. As in the single class case, we observe that the the average blocking probabilities under the MV(2) scheme are significantly smaller than that under the randomized static scheme.

We now numerically verify the asymptotic insensitivity of the system under the MV(d) scheme as was indicated in Remark 17. To do so, we compare the average blocking probability of jobs under the MV(d) scheme for different service time distributions. We set the following parameter values: $L = 1$, $M = 2$, $d = 2$, $N = 100$, $\gamma_1 = \gamma_2 = 0.5$, $C_1 = 20$, and $C_2 = 25$. In Table 2.2, average blocking probability is shown for different values of λ and for the following distributions.

1. *Constant*: We consider job length distribution having the cumulative distribution given by $F(x) = 0$ for $0 \leq x < 1$, and $F(x) = 1$, otherwise.

[2]The 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.

Table 3.2: Average blocking probability of jobs in the multi-class case

λ	MV(2) Scheme		Randomized Static Scheme	
	Class 1	Class 2	Class 1	Class 2
2.0	0.0077	0.0002	0.0969	0.0245
3.0	0.0740	0.0029	0.2047	0.0592
5.0	0.3968	0.0386	0.4349	0.1589
10.0	0.6054	0.2747	0.7728	0.3787
15.0	0.7515	0.4547	0.8967	0.5148

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 mean service time is 1. We see from Table 3.3 that the change in the average blocking probability is insignificant when the service time distribution is changed keeping the same mean. This supports the fact that under the MV(d) scheme the system becomes insensitive to service time distributions as $N \to \infty$.

Table 3.3: Asymptotic insensitivity of the MV(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

3.6 Conclusion

In this chapter, we characterized the performance of randomized dynamic job assignment schemes in reducing the average blocking probability of jobs arriving at a multi-class heterogeneous Erlang loss system. In particular, we considered the MV(d) scheme, where each incoming job request is assigned to the server having the maximum amount of unused resource among a set of d servers, sampled uniformly at random from the entire system. The system operating under the MV(d) scheme was analyzed in the large system limit. It was shown that in the large system limit the evolution of the empirical distribution of states of the servers can be described by a mean field limit. We established the existence of a unique equilibrium point of the mean field and showed that it characterizes the stationary occupancy distribution of the servers. We also established independence of the servers in the limiting system. Numerical studies revealed that the average blocking probability of jobs in the system under the MV(d) scheme is significantly lower than that under the randomized static scheme, where jobs are assigned to the servers independently of server states. The MV(d) scheme was further shown to be nearly optimal in terms of reducing the average blocking probability of jobs. Numerical evidence supporting asymptotic insensitivity of the system under the MV(d) scheme was also provided. The results presented in this chapter also appeared in [62, 63].

Part II

Opinion Dynamics in Large Social Networks

Chapter 4

Binary Opinion dynamics with Biased and Stubborn Agents

In this chapter, we investigate the impact of random interactions between agents in a social network on the diffusion of opinions in the network. The agents are assumed to be mobile and hence do not have any fixed set of neighbors. We consider the case where the opinion of each agent is a binary variable taking values in the set $\{0, 1\}$. This represents scenarios where every agent in the network has to choose from two available options. Each agent updates its opinion at random instants and interacts with other agents while performing the update. We consider two rules of interaction differing in the number individuals contacted during the update: (1) the *voter rule* in which a single agent is contacted during an update; (2) the *majority rule*, in which multiple agents are contacted during an update. Under each rule, we consider two different scenarios: (1) where the agents are 'biased' towards one of the opinions, (2) where different agents have different propensities to change their opinions. For each scenario, we characterize the equilibrium state of the network as a function of the initial opinions of the agents and the number of agents in the network. We also characterize the time to reach the equilibrium state as a function of the aforementioned factors.

4.1 Introduction

With the widespread use of online social networking, opinions of individuals are constantly being shaped by social interactions. Understanding how individual opinions are affected by social interactions and what global opinion structure emerges from such interactions are important in many contexts such as economics, politics and psychology. The dynamics of opinion formation in a network is also important in the context of viral marketing and can provide useful guidelines for targeted message delivery to users in a network. Consequently, modeling social interactions both analytically and through empirical studies has gained vast attention in different fields of study.

Mathematical models of social interaction treat opinion of each individual in a social network as a variable taking values in either a discrete or a continuous subset of the Euclidean space. Although this may seem too reductive to capture the complexity of choices made by real individuals, in everyday situations, individuals in a network are often faced with only a limited number of choices (often as few as two) concerning a specific issue, e.g., pro-/anti-government, Windows/Linux, Democrat/Republican, etc. Thus, a vast body of literature treats opinions of individuals as binary variables taking values in the set $\{0, 1\}$.

The interactions between agents in a social network are generally modeled using simple rules that capture the essential features of individuals in a society such as their tendency to mimic their neighbors or to conform with the majority opinion in local neighborhoods. One of the models, extensively analyzed in the literature, is the *voter model* [5, 6, 7] or the *voter rule*, where an agent randomly samples one of its neighbors at an instant when it decides to update its opinion. The updating agent then adopts the opinion of the sampled neighbor. This simple rule captures the tendency of an individual to mimic other individuals in the society. Because of its simplicity the rule has been analyzed under a variety of network topologies [64] that assume connectedness of the underlying graph.[1] It is known that, under the voter rule, any connected network converges to a *consensus*, where all individuals adopt the same opinion. It is of interest to determine the probability with which consensus is reached on a specific opinion and the time it takes for the network to converge to the consensus state.

Another rule studied in this context is the *majority rule* model [9, 65, 66]. In it, instead of sampling a single individual, an updating agent consults multiple individuals while performing the update and adopts the choice of the majority of the sampled neighbors. This rule captures the tendency of the individuals to conform with the majority opinion in their local neighborhoods. In a fully connected network, the majority rule also leads to a consensus among agents. However, the rate at which consensus is reached is faster than that under the voter rule.

In all the prior works on the voter models and the majority rule models, it is assumed

[1]Connectedness implies that every individual is connected to every other individual either directly or via immediate neighbors.

that an agent's decision to update its opinion does not depend on the current opinion of the agent. It is also assumed that all agents in the network have the same propensity to change their opinions. However, in a real scenario an agent may be 'biased' towards a specific opinion in the sense that if it holds its 'preferred' opinion, then the probability with which it updates its opinion is small. We may also encounter situations where some of the agents update their choices less frequently than others (irrespective of their current opinions). In this chapter, we focus on these two scenarios.

4.1.1 Related literature

There is a rich and growing literature that studies diffusion of technologies and opinions in large social networks. Early studies on this topic [67] considered models for growth of new products in a society. Such models are similar to the *susceptible-infected-susceptible* (SIS) and *susceptible-infected-removed* (SIR) models which describe the spread of epidemics or infections in networks [68, 69]. In both the SIS and SIR models, diffusion occurs through infected agents, infecting their susceptible neighbors. The process of recovery from the infected state to the susceptible state, however, is assumed to be independent of the number of susceptible agents in the neighborhood. This is different from the models considered in this chapter, where all transitions depend on the configuration of the local neighborhood of the agents.

An interesting line of research evolved in the 1990's where opinion dynamics was treated in the Bayesian setting [70, 71, 72, 73, 74, 75]. In this setting, individuals observe the actions of other individuals in the society and update their beliefs about an underlying state variable. This process is referred to as *social learning*. The primary question of interest in these models is whether social learning leads to a society where individuals adopt the technologies that produce higher payoffs for them. Although, in this chapter, we are interested in similar questions, our models are different from the models of social learning since we do *not* assume that the updating agents can observe the opinions of all the other agents in the network.

Another line of research considers opinion dynamics models in the non-Bayesian setting. The voter models and the majority rule models fall under this category. Such models were first studied by DeGroot [76] where the agents were assumed to update their opinions (assumed to be continuous variables within a certain range) synchronously by averaging the opinions of their neighbors. This is equivalent to the synchronous average consensus algorithms considered in [77] and thus can be analyzed using similar techniques. The 'voter model' with binary opinions was first studied independently in [5] and [6]. It was assumed that an agent simply copies the opinion of a randomly sampled neighbor at an instant of update. Due to its simplicity, the voter model soon became popular and was analyzed under a variety of network topologies, e.g., finite integer lattices in different dimensions [7, 78], heterogeneous graphs [79], Erdos-Renyi random graphs and random geometric graphs [64] etc. In [80, 81], the voter model was studied under the presence of stubborn individuals who do not update their opinions. In such a scenario, the network cannot reach a consensus because of the presence of stubborn agents having both opinions. Using coalescing random walk techniques the average opinion in the network and the variance of opinions were computed at steady state.

The majority rule model was first introduced in [8], where it was assumed that, at every iteration, groups of random sizes are formed by the agents. Within each group, the majority opinion is adopted by all the agents. Under this rule, it was shown that consensus is achieved on a particular opinion with high probability only if the initial fraction of agents having that opinion is more than a certain critical value. Furthermore, the time to reach consensus was shown to scale as logarithm of the network size (number of agents). Similar models with fixed (odd) group size were considered in [9, 65]. It was shown that for finite dimensional integer lattices the consensus time grows as a power law in the number of agents in the network.

A deterministic version of the majority rule model, where an agent, instead of randomly sampling some of its neighbors, adopts the majority opinion among all its neighbors, is considered in [82, 83, 84, 85]. In such models, given the graph structure of the network, the opinions of the agents at any time is a deterministic function of the initial opinions of the agents. The interest there is to find out the the initial distribution of opinions for which the network converges to some specific absorbing state. In social networks, where the neighborhood of each agent is large, such majority rule dynamics involves complex computation by each updating agent at each update instant. Our interest in this chapter is on scenarios where the agents are mobile and do not have any fixed neighborhoods. We therefore consider a randomized version of the majority rule.

4.1.2 Contributions

In this chapter, we study binary opinion dynamics under the voter model and the majority rule model. Under each model, we consider two different scenarios. In the first scenario, the agents are assumed to be 'biased' towards a specific opinion. More specifically, the agents having one of the two opinions (the 'preferred' opinion) are assumed to update their choices less frequently than agents having the other opinion. In the second scenario, different agents are assumed to have different propensities to change their opinions, irrespective of their current opinions.

For the voter model with biased agents, we derive a closed form expression of the probability with which consensus is reached on the 'preferred' opinion. It is observed that this probability rapidly increases to 1 as the number of agents in the network grows. This is unlike the case with unbiased agents, where the probability to reach consensus on a particular opinion remains constant for all network sizes. Using mean field techniques, we derive an estimate of the average time taken for the network to reach consensus. It is observed that the mean consensus time grows as logarithm of the network size. This is in contrast to the case with unbiased agents, where the mean consensus time grows linearly with the number of agents.

Next we analyze the voter model under a scenario where the agents are categorized into two groups. Agents belonging to the first group are assumed to update their opinions with a lesser (but non-zero) probability than the agents belonging to the second group. Closed form expression of the probability with which the network reaches consensus on a particular opinion is derived. It is observed that this probability does not depend on the number of agents in the network. Furthermore, an approximate expression of the mean consensus time is derived for large network sizes. It is found that the mean consensus time grows linearly with the network size.

For the majority rule model with biased agents, we derive a closed form expression for the probability with which consensus is achieved on the preferred opinion. It is observed that, unlike the voter model, consensus is achieved on the preferred opinion (with high probability) only if the initial fraction of agents having that opinion is above a certain threshold. This threshold is determined from the mean field analysis of the model. An estimate of the mean consensus time is also found from the mean field model. It suggests that the mean consensus time grows as logarithm of the number of agents in the network.

Finally, we consider the majority rule model when there are 'stubborn' agents in the network. The stubborn agents are assumed to have fixed opinions at all times. Therefore, in this case consensus can never be reached. We analyze the equilibrium distribution of opinions among the non-stubborn agents using mean field techniques. Depending on the system parameters, the mean field is shown to have either multiple stable equilibrium points or a unique stable equilibrium point within the range of interest. As the system size grows, the equilibrium distribution of opinions among non-stubborn agents is shown to converge to a mixture of Dirac measures concentrated on the equilibrium points of the mean field. This suggests a *metastable* behavior of the system where the system moves back and forth between stable configurations, spending long intervals in each configuration. The conditions for metastability are obtained in terms of the system parameters.

4.1.3 Organization

The rest of the chapter is organized as follows. In Section 4.2, we introduce the voter model. In Subsections 4.2.1 and 4.2.2, we analyze the voter model with 'biased' agents and 'stubborn' agents, respectively. Section 4.3 introduces the majority rule model. In Subsections 4.3.1 and 4.3.2, we analyze the majority rule model with 'biased' and 'stubborn' agents, respectively. Finally, the chapter is concluded in Section 4.4.

4.2 The voter models

Let us consider a network consisting of N social agents, where each agent can communicate with every other agent. This represents a scenario where the agents are mobile and therefore do not have fixed sets of neighbors. Opinion of each agent is assumed to be a binary variable taking values in the set $\{0, 1\}$. Initially, every agent adopts one of the two opinions. The agents then consider updating their opinions at points of independent unit rate Poisson processes associated with themselves. At a point of the Poisson process associated with itself, an agent either updates its opinion or retains its past opinion. In case the agent decides to update its opinion, it samples an agent uniformly at random (with replacement) from the network[2] and adopts the opinion of the sampled agent. The agent sampled by the updating agent can be seen as the neighbor of the updating agent at the update instant.

Below we consider two different scenarios: (1) where the agents are 'biased towards a specific opinion, and (2) where the agents have different propensities to change their past opinions.

4.2.1 The voter model with biased agents

We first consider the case where the agents are 'biased' towards one of the two opinions. Without loss of generality, we assume that all agents in the network prefer opinion $\{1\}$ to opinion $\{0\}$. This is modeled as follows: Each agent with opinion $i \in \{0, 1\}$ updates its opinion at a point of the unit rate Poisson process associated with itself with probability q_i and retains its opinion with probability $p_i = 1 - q_i$. This is equivalent to an agent with

[2] In the large N limit sampling with or without replacement does not make any difference.

opinion i updating its opinion at all points of a Poisson process with rate q_i. In case the agent decides to update its opinion, the update occurs following the voter rule discussed in the beginning of this section. We assume $q_0 > q_1$ ($p_1 > p_0$) to imply that the agents having opinion $\{0\}$ update their opinions more frequently than the agents having opinion $\{1\}$. In the above sense, the agents are biased towards opinion $\{1\}$.

Clearly, in this case, the network gets absorbed (in a finite time) in a state where all the agents adopt the same opinion. This is referred to as the *consensus state*. Our interest is to find out the probability with which consensus is achieved on the preferred opinion $\{1\}$. This is known as the *exit probability* of the network. We also intend to characterize the mean time to reach the consensus state.

The case $q_1 = q_0 = 1$ is referred to as the voter model with unbiased agents, which has been analyzed in [5, 6]. It is known that for unbiased agents the probability with which consensus is reached on a particular opinion is simply equal to the initial fraction α of agents having that opinion and the expected time to reach consensus for large N is approximately given by $Nh(\alpha)$, where $h(\alpha) = -[\alpha \ln(\alpha) + (1-\alpha)\ln(1-\alpha)]$. We now proceed to characterize these quantities for the voter model with biased agents.

Let $X^{(N)}(t)$ denote the number of agents with opinion $\{1\}$ at time $t \geq 0$. Clearly, $X^{(N)}(\cdot)$ is a Markov process on state space $\{0, 1, \ldots, N\}$, with possible jumps at the points of a rate N Poisson process. This rate N process is referred to as the *global clock*. All states, except the states 0 and N, form an open communicating class; the states 0 and N are the absorbing states. Therefore, with probability 1, the process gets absorbed in one of the absorbing states in a finite time. Our interest is to find the probability with which the process $X^{(N)}(\cdot)$ gets absorbed in state N.

Proposition 4.2.1. *The probability $E_N(n)$ with which the process $X^{(N)}(\cdot)$ gets absorbed in state N starting with state $n \in \{1, 2 \ldots, N\}$ is given by*

$$E_N(n) = \frac{1 - r^n}{1 - r^N}, \tag{4.1}$$

where $r = q_1/q_0 < 1$ and $E_N(0) = 0$.

Proof. Given that the process $X^{(N)}(\cdot)$ is in state k at one point of the global clock, it transits to state $k+1$ at the next point of the global clock only if one of the agents having opinion $\{0\}$ updates its opinion to opinion $\{1\}$. The probability with which any one of the $N-k$ agents having opinion $\{0\}$ decides to update its opinion is given by $q_0 \times (N-k)/N$. The probability with which the updating agent samples an agent with opinion $\{1\}$ is given by k/N. Hence, the total probability with which the process $X^{(N)}(\cdot)$ transits to from the state k to the state $k+1$ is given by

$$p(k \to k+1) = \frac{k(N-k)}{N^2} q_0. \tag{4.2}$$

Similarly, the probability of transition from the state k to the state $k-1$ is given by

$$p(k \to k-1) = \frac{k(N-k)}{N^2} q_1. \tag{4.3}$$

Therefore, the probability with which no transition occurs between two consecutive points of the global clock is

$$p(k \to k) = 1 - p(k \to k+1) - p(k \to k-1). \tag{4.4}$$

We now use these transition probabilities to derive the expression of the exit probability $E_N(n)$ as follows. Since $X^{(N)}(\cdot)$ is Markov, $E_N(n)$ must satisfy the following recursive relationship

$$E_N(n) = p(n \to n+1)E_N(n+1) + p(n \to n-1)E_N(n-1) + p(n \to n)E_N(n). \tag{4.5}$$

Using (4.2), (4.3), and (4.4) it is easy to see that (4.5) simplifies to

$$E_N(n+1) - E_N(n) = r(E_N(n+1) - E_N(n)), \tag{4.6}$$

where $r = q_1/q_0$. This homogeneous, linear recursive relationship can be solved easily using boundary conditions $E_N(0) = 0$ and $E_N(N) = 1$. The solution is given by (4.1).[3] □

In terms of the initial fraction $\alpha = n/N$ of agents having opinion $\{1\}$, (4.1) can be rewritten as

$$E_N(\alpha) = \frac{1 - r^{N\alpha}}{1 - r^N}, \tag{4.7}$$

where $E_N(\alpha)$ denotes the probability with the process $X^{(N)}(\cdot)$ is absorbed in state N starting with α fraction of agents having opinion $\{1\}$.

[3] For $q_0 = q_1$, we have $E_N(n) = n/N$, which coincides with the known results.

Clearly, for $q_0 > q_1$, we have $r < 1$. Hence, from (4.7) it is clear that as N increases the exit probability rapidly increases to 1 for all α. This is in contrast to the case with unbiased agents ($q_0 = q_1 = 1$) where the exit probability remains constant at α for all values of N.

We now characterize the mean time $\bar{T}_N(\alpha)$ to reach the consensus state starting from α fraction of agents having opinion $\{1\}$. To do so, we consider the empirical measure process $x^{(N)}(\cdot) = X^{(N)}(\cdot)/N$. The process $x^{(N)}(\cdot)$ jumps from the state x to the state $x+1/N$ when one of the $N(1-x)$ agents having opinion $\{0\}$ updates (with probability q_0) its opinion by interacting with an agent with opinion $\{1\}$. Since the agents update their opinions at points of independent unit rate Poisson processes, the rate at which one of the $N(1-x)$ agents having opinion $\{0\}$ decides to update its opinion is $N(1-x)q_0$. The probability with which the updating agent interacts with an agent with opinion $\{1\}$ is x. Hence, the total rate of transition from x to $x+1/N$ is given by

$$r(x \to x + 1/N) = q_0 N x(1-x). \tag{4.8}$$

Similarly, the rate of transition from x to $x - 1/N$ is given by

$$r(x \to x - 1/N) = q_1 N x(1-x). \tag{4.9}$$

From the above transition rates it can be easily seen that the generator of the process $x^{(N)}(\cdot)$ converges uniformly as $N \to \infty$ to the generator of the process $x(\cdot)$ which satisfies the following differential equation

$$\dot{x}(t) = (q_0 - q_1)x(t)(1 - x(t)). \tag{4.10}$$

Thus, if $x^{(N)}(0) \Rightarrow x(0)$ as $N \to \infty$, then by Theorem C.2.2 and Theorem C.2.3, $x^{(N)}(\cdot) \Rightarrow x(\cdot)$ as $N \to \infty$. In other words, $x(\cdot)$ is the mean field limit of the system.

Since $q_0 > q_1$ and $x(t) \in [0, 1]$ for all $t \geq 0$, we have from (4.10) that $\dot{x}(t) \geq 0$ for all $t \geq 0$. Hence, $x(t) \to 1$ as $t \to \infty$. The mean consensus time $\bar{T}_N(\alpha)$ for large N can therefore be approximated by the time taken by the process $x(t)$ to reach the state $1-1/N$ (which corresponds to the situation where all the agents except one agent have opinion $\{1\}$) starting with $x(0) = \alpha$. Solving (4.10) we have

$$\ln\left(\frac{x(t)(1-\alpha)}{(1-x(t))\alpha}\right) = (q_0 - q_1)t. \tag{4.11}$$

Now using the limits discussed above, the mean consensus time $\bar{T}_N(\alpha)$ can be found as

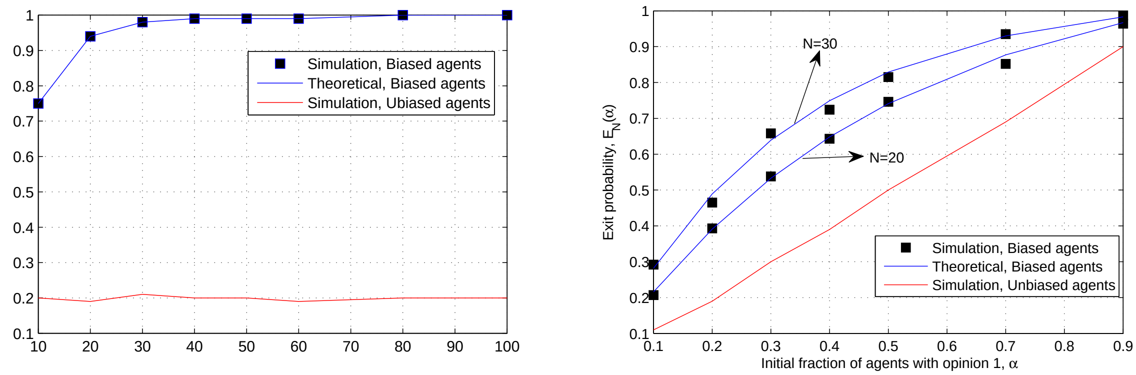

(a) Exit probability $E_N(\alpha)$ as a function of the number of agents N. Parameters: $q_0 = 1$, $q_1 = 0.5$, $\alpha = 0.2$.

(b) Exit probability $E_N(\alpha)$ as a function of the initial fraction α of agents with opinion $\{1\}$. Prameters: $q_0 = 1, q_1 = 0.9, N = 100$.

Figure 4.1: Exit probability under the voter model with biased agents

$$\bar{T}_N(\alpha) = \frac{1}{q_0 - q_1} \ln(N-1) - \frac{1}{q_0 - q_1} \ln\left(\frac{\alpha}{1-\alpha}\right) = O\left(\frac{1}{|q_0 - q_1|} \ln(N-1)\right) \tag{4.12}$$

Clearly, the mean consensus time scales as $O(\ln N)$. This is in contrast to the voter model with unbiased agents where the mean consensus time is known to increase linearly with the network size N. Thus, in the case with biased agents, the network reaches the consensus state exponentially faster than that in the case with unbiased agents.

Numerical Results: We now numerically study the exit probability $E_N(\alpha)$ and the mean consensus time $\bar{T}_N(\alpha)$ as functions of the network size N and the initial fraction α of agents having opinion $\{1\}$.

In Figure 4.1(a), we plot the exit probability for both biased ($q_0 > q_1$) and unbiased ($q_0 = q_1 = 1$) cases as functions of the number of agents N for $\alpha = 0.2$. For the biased case, we have chosen $q_0 = 1, q_1 = 0.5$. We observe that in the biased case the exit probability rapidly increases to 1 with the increase N. This is in contrast to the unbiased case, where the exit probability remains constant at α for all N.

In Figure 4.1(b), we plot the exit probability $E_N(\alpha)$ as a function of α for both biased and unbiased cases. We observe that, in the biased case, the exit probability increases to 1 at a faster rate than that in the unbiased case. We also observe that, in the unbiased case, exit probability grows linearly with α, independent of N.

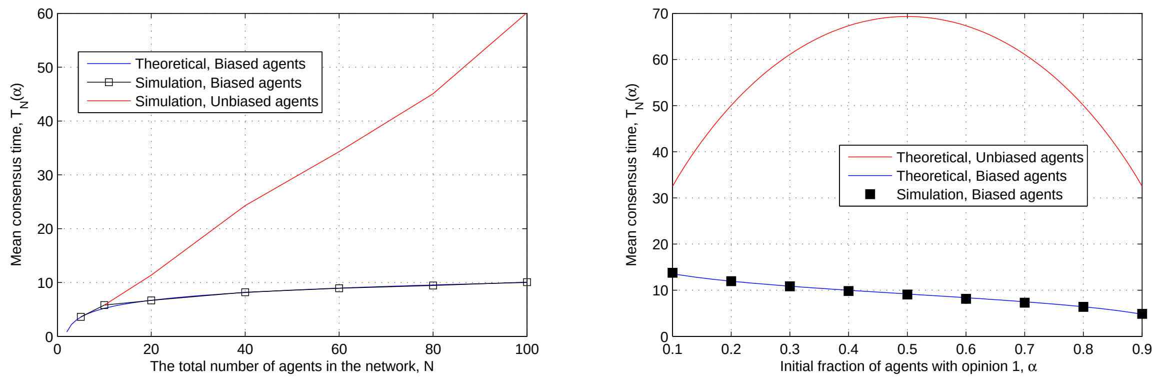

(a) Mean consensus time $\bar{T}_N(\alpha)$ as a function of the number of agents N. Parameters: $q_0 = 1$, $q_1 = 0.5$, $\alpha = 0.4$.

(b) Mean consensus time $\bar{T}_N(\alpha)$ as a function of the initial fraction α of agents with opinion $\{1\}$. Parameters: $q_0 = 1, q_1 = 0.6, N = 200$.

Figure 4.2: Mean consensus time under the voter model with biased agents

In Figure 4.2(a), we plot the mean consensus time $\bar{T}_N(\alpha)$ for both the biased and unbiased cases as functions of N for $\alpha = 0.4$. We observe that, in the biased case, the consensus state is reached in a time exponentially smaller than that in the unbiased case. This is because the bias of the agents towards one of the opinions drives the system to consensus much faster.

In Figure 4.2(b), we plot the mean consensus time as a function of α for both biased an unbiased cases. The network size is kept fixed at $N = 100$. We observe that for the unbiased case, the consensus time increases for $\alpha \in (0, 0.5)$ and decreases for $\alpha \in (0.5, 1)$. In contrast, for the biased case, the consensus time steadily decreases with the increase in α. This is expected since, in the unbiased case, consensus is achieved faster on a particular opinion if the initial number agents having that opinion is more than the initial number of agents having the other opinion. On the other hand, in the biased case, consensus is achieved with high probability on the preferred opinion and therefore increasing the initial fraction of agents having the preferred opinion always decreases the mean consensus time.

4.2.2 The voter model with stubborn agents

We now consider the case where different agents have different propensities to change their opinions. This is modeled as follows: Each agent in the network is assumed to belong to one of the two disjoint sets $\mathcal{S}$ and $\mathcal{R}$. We denote by γ_S and $\gamma_R = 1 - \gamma_S$ the fractions of agents belonging to the sets $\mathcal{S}$ and $\mathcal{R}$, respectively. Each agent belonging to the set $\mathcal{S}$ ($\mathcal{R}$) updates its opinion with probability q_S (q_R) at a point of the unit rate Poisson process associated with itself and retains its opinion with probability $p_S = 1 - q_S$ ($p_R = 1 - q_R$). The updates occur according to the voter rule, discussed in the beginning of this section. The probabilities q_S and q_R determine the degrees of 'stubbornness' of agents belonging to the sets $\mathcal{S}$ and $\mathcal{R}$, respectively. We set $q_S < q_R$ to imply that the agents belonging to the set $\mathcal{S}$ update their opinions less frequently than the agents belonging to the set $\mathcal{R}$.

We note that, unlike the voter model with biased agents, in this model, the probability with which an agent updates its opinion does not depend on the current opinion of the agent and is determined by the set ($\mathcal{S}$ or $\mathcal{R}$) the agent belongs to. We also point out the difference between the model discussed above and the models considered in [81, 80]. In [81, 80], it was assumed that the 'stubborn' agents do not change their opinions at any time. This implies that a consensus can never be achieved in such cases due to the presence of 'stubborn' agents having both opinions at all times. However, in our model, agents belonging to both the sets $\mathcal{S}$ and $\mathcal{R}$ are assumed to have non-zero update probabilities (i.e, $q_S, q_R \neq 0$). Hence, in our model, consensus is always reached in a finite time (with probability 1). Below we characterize the exit probability and mean time to reach consensus for the model under consideration.

The evolution of the network can be described by a two dimensional Markov process $X^{(N)}(\cdot) = (X_S^{(N)}(\cdot), X_R^{(N)}(\cdot))$, where $X_S^{(N)}(t)$ and $X_R^{(N)}(t)$ denote the numbers of agents with opinion $\{1\}$ in sets $\mathcal{S}$ and $\mathcal{R}$, respectively, at time t. Let (m, n) be the state of the process at some instant. The process transits to state $(m+1, n)$ when one of the $N\gamma_S - m$ agents with opinion $\{0\}$ in the set $\mathcal{S}$ updates its opinion by interacting with an agent with opinion $\{1\}$. The rate at which any one of the $N\gamma_S - m$ agents having opinion $\{0\}$ in set $\mathcal{S}$ decides to update its opinion is $(N\gamma_S - m)q_S$. The probability with which the updating agent samples an agent with opinion $\{1\}$ from the entire network is $(m+n)/N$. Hence, the total rate of this transition is given by

$$r\left((m, n) \to (m+1, n)\right) = \frac{(N\gamma_S - m)(m+n)}{N} q_S \tag{4.13}$$

Similarly, the rates of other possible transitions are given by

$$r\left((m,n)\to(m,n+1)\right)=\frac{(N\gamma_R-n)(m+n)}{N}q_R \tag{4.14}$$

$$r\left((m,n)\to(m-1,n)\right)=\frac{m(N-m-n)}{N}q_S \tag{4.15}$$

$$r\left((m,n)\to(m,n-1)\right)=\frac{n(N-m-n)}{N}q_R \tag{4.16}$$

Proposition 4.2.2. *Let $E_N(\alpha_S,\alpha_R)$ denote the probability with which the network with N agents reaches a consensus state with all agents having opinion $\{1\}$ starting with α_S (resp. α_R) fraction of agents of the set $\mathcal{S}$ (resp. $\mathcal{R}$) having opinion $\{1\}$. Then*

$$E_N(\alpha_S,\alpha_R)=\frac{q_R\gamma_S\alpha_S+q_S\gamma_R\alpha_R}{q_R\gamma_S+q_S\gamma_R}. \tag{4.17}$$

Proof. Let $\mathcal{F}_t=\sigma(X^{(N)}(s),0\le s\le t)$ denote the history of the process $X^{(N)}(\cdot)$ upto time $t\ge 0$. Consider the process $X_S^{(N)}(\cdot)/q_S+X_R^{(N)}(\cdot)/q_R$. The conditional drift of the process from time t to time $t+h$ is given by

$$\begin{aligned}
&\mathbb{E}\left[\frac{X_S^{(N)}(t+h)}{q_S}+\frac{X_R^{(N)}(t+h)}{q_R}-\frac{X_S^{(N)}(t)}{q_S}-\frac{X_R^{(N)}(t)}{q_R}|\mathcal{F}_t\right]\\
&=\left(\left(\frac{1}{q_S}\right)r\left((m,n)\to(m+1,n)\right)+\left(\frac{1}{q_R}\right)r\left((m,n)\to(m,n+1)\right)\right.\\
&\left.-\left(\frac{1}{q_S}\right)r\left((m,n)\to(m-1,n)\right)-\left(\frac{1}{q_R}\right)r\left((m,n)\to(m,n-1)\right)\right)h+o(h)\\
&=o(h)
\end{aligned}$$

The third equality follows from the second since the coefficient of h vanishes using (4.13), (4.14), (4.15), and (4.16). Thus, the process $X_S^{(N)}(\cdot)/q_S+X_R^{(N)}(\cdot)/q_R$ is an $\mathcal{F}_t$ martingale. Let T denote the random time the process $X^{(N)}(\cdot)$ hits the consensus state. Clearly, T is an $\mathcal{F}_t$ stopping time. Hence, using optional sampling theorem we have

$$\mathbb{E}\left[\frac{X_S^{(N)}(T)}{q_S}+\frac{X_R^{(N)}(T)}{q_R}\right]=\mathbb{E}\left[\frac{X_S^{(N)}(0)}{q_S}+\frac{X_R^{(N)}(0)}{q_R}\right]=\frac{N\gamma_S\alpha_S}{q_S}+\frac{N\gamma_R\alpha_R}{q_R} \tag{4.18}$$

The left hand side of the above equation can be written as

$$\mathbb{E}\left[\frac{X_S^{(N)}(T)}{q_S} + \frac{X_R^{(N)}(T)}{q_R}\right] = \left(\frac{N\gamma_S}{q_S} + \frac{N\gamma_R}{q_R}\right) E_N(\alpha_S, \alpha_R) + 0 \times (1 - E_N(\alpha_S, \alpha_R)) \quad (4.19)$$

Hence we obtain

$$E_N(\alpha_S, \alpha_R) = \frac{\frac{N\gamma_S\alpha_S}{q_S} + \frac{N\gamma_R\alpha_R}{q_R}}{\frac{N\gamma_S}{q_S} + \frac{N\gamma_R}{q_R}} \quad (4.20)$$

which simplifies to (4.17). □

Remark 18. From (4.17) we see that the exit probability does not depend on the number of agents N and is a function of the initial fractions α_S, α_R; the probabilities q_S, q_R which determine the stubbornness of the agents; and the proportions γ_S, γ_R which define the sizes of the sets $\mathcal{S}$ and $\mathcal{R}$, respectively. We also observe that for $\alpha_S = \alpha_R = \alpha$, the exit probability is given by $E_N(\alpha, \alpha) = \alpha$, which is independent of q_S and q_R.

The mean time $\bar{T}_N(\alpha_S, \alpha_R)$ to reach consensus starting with α_S (resp. α_R) fraction of agents of the set $\mathcal{S}$ (resp. $\mathcal{R}$) having opinion $\{1\}$ can be computed using the first step analysis of the empirical measure process $x^{(N)}(\cdot) = (X_S^{(N)}(\cdot)/N\gamma_S, X_R^{(N)}(\cdot)/N\gamma_R)$. The process $x^{(N)}(\cdot)$ changes its state only at points of a rate N Poisson point process, referred to as the global clock. The probability $p((x, y) \to (x + 1/N\gamma_S, y))$ with which the process transits from the state (x, y) at one point of the global clock to the state $(x + 1/N\gamma_S, y)$ at the next point of the global clock is given by

$$p\left((x, y) \to \left(x + \frac{1}{N\gamma_S}, y\right)\right) = \gamma_S(1 - x)(\gamma_S x + \gamma_R y)q_S. \quad (4.21)$$

Similarly, the probabilities for the other possible transitions are given by

$$p\left((x,y)\to\left(x,y+\frac{1}{N\gamma_R}\right)\right)=\gamma_R(1-y)(\gamma_S x+\gamma_R y)q_R \tag{4.22}$$

$$p\left((x,y)\to\left(x-\frac{1}{N\gamma_S},y\right)\right)=\gamma_S x(1-\gamma_S x-\gamma_R y)q_S \tag{4.23}$$

$$p\left((x,y)\to\left(x,y-\frac{1}{N\gamma_R}\right)\right)=\gamma_S y(1-\gamma_S x-\gamma_R y)q_R \tag{4.24}$$

$$\begin{aligned}p\left((x,y)\to(x,y)\right)&=1-p\left((x,y)\to\left(x+\frac{1}{N\gamma_S},y\right)\right)-p\left((x,y)\to\left(x,y+\frac{1}{N\gamma_R}\right)\right)\\&-p\left((x,y)\to\left(x-\frac{1}{N\gamma_S},y\right)\right)-p\left((x,y)\to\left(x,y-\frac{1}{N\gamma_R}\right)\right)\end{aligned} \tag{4.25}$$

Since the process $x^{(N)}(\cdot)$ is Markov and the average gap between two points of the global clock is $1/N$, we have the following recursive relation

$$\begin{aligned}\bar{T}_N(x,y)=p\left((x,y)\to\left(x+\frac{1}{N\gamma_S},y\right)\right)\left(\bar{T}_N\left(x+\frac{1}{N\gamma_S},y\right)+\frac{1}{N}\right)\\+p\left((x,y)\to\left(x,y+\frac{1}{N\gamma_R}\right)\right)\left(\bar{T}_N\left(x,y+\frac{1}{N\gamma_R}\right)+\frac{1}{N}\right)\\+p\left((x,y)\to\left(x-\frac{1}{N\gamma_S},y\right)\right)\left(\bar{T}_N\left(x-\frac{1}{N\gamma_S},y\right)+\frac{1}{N}\right)\\+p\left((x,y)\to\left(x,y-\frac{1}{N\gamma_R}\right)\right)\left(\bar{T}_N\left(x,y-\frac{1}{N\gamma_R}\right)+\frac{1}{N}\right)\\+p\left((x,y)\to(x,y)\right)\left(\bar{T}_N(x,y)+\frac{1}{N}\right)\end{aligned} \tag{4.26}$$

Now using (4.21), (4.22), (4.23), (4.24) and the Taylor series expansion of $\bar{T}_N(\cdot,\cdot)$ of second order around the point (x,y) we have that for large N

$$\begin{aligned}&\gamma_R q_S(y-x)\frac{\partial\bar{T}(x,y)}{\partial x}+\frac{q_S((\gamma_S+1)x+\gamma_R y-2x(\gamma_S x+\gamma_R y))}{2N\gamma_S}\frac{\partial^2\bar{T}(x,y)}{\partial x^2}\\&+\gamma_S q_P(x-y)\frac{\partial\bar{T}(x,y)}{\partial y}+\frac{q_R(\gamma_S x+(\gamma_R+1)y-2y(\gamma_S x+\gamma_R y))}{2N\gamma_R}\frac{\partial^2\bar{T}(x,y)}{\partial y^2}=-1\end{aligned} \tag{4.27}$$

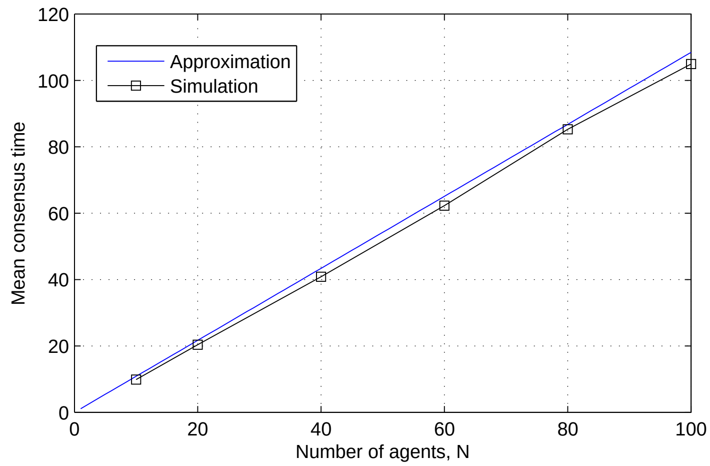

Figure 4.3: Consensus time as a function of the network size N. Parameters: $q_S = 0.3$, $q_R = 1$, $\alpha_S = \alpha_R = 0.8$, $\gamma_S = \gamma_R = 0.5$.

with boundary condition $\bar{T}_N(0,0) = \bar{T}_N(1,1) = 0$. An approximate solution of the above partial differential equation is given as by

$$\bar{T}_N(x,y) = N\left(\frac{\gamma_S}{q_S} + \frac{\gamma_R}{q_R}\right) h\left(\frac{\gamma_S q_R}{\gamma_S q_R + \gamma_R q_S} x + \frac{\gamma_R q_S}{\gamma_S q_R + \gamma_R q_S} y\right), \tag{4.28}$$

where $h(z) = -(z \ln z + (1-z)\ln(1-z))$. The approximation is obtained by noting that the above solution is exact for the cases $\gamma_S = 1, \gamma_R = 0$ and $\gamma_S = 0, \gamma_R = 1$. Moreover, putting the solution in (4.27) we see that the terms containing first order partial derivatives vanish and the terms containing the second order partial derivatives simplify approximately to -1.

Numerical results: To numerically investigate how consensus time varies with the system size N, in Figure 4.3 we plot the mean consensus time of 1000 independent runs of a network with the following parameters: $q_S = 0.3$, $q_R = 1$, $\alpha_S = \alpha_R = 0.8$, $\gamma_S = \gamma_R = 0.5$. We observe that the mean consensus time grows linearly with N. In the figure, we have also plotted the mean consensus time obtained using (4.28). We observe a close match between the simulation result and the approximate result which suggests that the approximation provided in (4.28) is accurate.

Next, we investigate how the consensus time varies with q_S the probability with which the stubborn agents update their opinions. To do so, we plot the mean consensus time as function of q_S for a system with the following parameters: $q_R = 1$, $\alpha_S = \alpha_R = 0.8$,

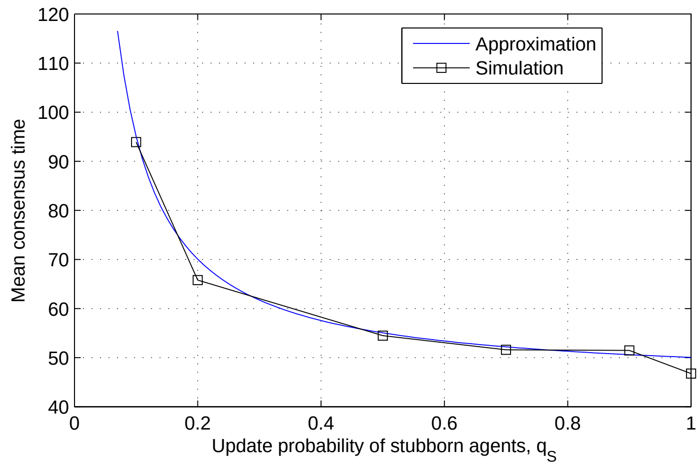

Figure 4.4: Consensus time as a function of the update probability q_S of stubborn agents. Parameters: $q_R = 1$, $\alpha_S = \alpha_R = 0.8$, $\gamma_S = 0.1 = 1 - \gamma_R$, $N = 100$.

$\gamma_S = 0.1 = 1 - \gamma_R$, $N = 100$. We see that the mean consensus time decreases with the increase in q_S. This is also suggested by the approximation (4.28).

4.3 The majority rule models

In this section, we consider models where an agent, instead of interacting with a single agent, interacts with multiple agents at an update instant. As before, we assume that the agents in the network consider updating their opinions at points of independent, unit rate Poisson point processes. At a point of the Poisson process associated with itself, an agent either retains its opinion or updates it. If the agent decides to update its opinion, then it interacts with $2K$ ($K \geq 1$) agents sampled uniformly at random (with replacement) from the network and adopts the opinion held by the majority of the $2K + 1$ agents which includes the $2K$ sampled agents and the updating agent itself.

As in the case of voter models, the decision of an agent to update its opinion is assumed to depend either i) on the current opinion held by the agent or ii) on the propensity of the agent to change opinions. Below we consider these two scenarios separately. To simplify analysis, we focus on the $K = 1$ case, where each agent interacts with two agents at its update instant. The case $K > 1$ is studied separately in Subsection 4.3.3.

We note that in the majority rule model discussed above, only one agent updates its state, at each time step, by interacting with a group of randomly sampled neighbors. This is different from the previous models studied in the literature [9, 65, 8], where all members of a group of interacting agents were assumed to update their opinions simultaneously.

4.3.1 The majority rule model with biased agents

As in Section 4.2.1, we first consider the case where the agents are 'biased' towards one of the two opinions. More specifically, we assume that an agent with opinion $i \in \{0, 1\}$ updates its opinion with probability q_i at a point of the unit rate Poisson process associated with itself. The agent retains its opinion with probability $p_i = 1 - q_i$. In case the agent decides to update its opinion, the update occurs according to the majority rule discussed in the beginning of this section. We assume $q_0 > q_1$ to imply that agents with opinion $\{0\}$ update their opinions more frequently than agents with opinion $\{1\}$. In the above sense, the agents are biased towards opinion $\{1\}$. The case $q_1 = q_0 = 1$ corresponds to the majority rule model with unbiased agents.

Let $X^{(N)}(t)$ denote the number of agents with opinion $\{1\}$ at time $t \geq 0$. Clearly, $X^{(N)}(\cdot)$ is a Markov process on state space $\{0, 1, \ldots, N\}$, with possible jumps at the points of a rate N Poisson process. We refer to the above rate N Poisson process as the *global clock* of the system. All states, except the states 0 and N, form a an open communicating class, and the states 0 and N are the absorbing states. Therefore, with probability 1, the process gets absorbed in one of the absorbing states in a finite time.

Proposition 4.3.1. *The probability $E_N(n)$ with which the process $X^{(N)}(\cdot)$ gets absorbed in state N starting from state $n \in \{1, 2, \ldots, N\}$ is given by*

$$E_N(n) = \frac{1}{(1+r)^{N-1}} \sum_{k=0}^{n-1} \binom{N-1}{k} r^k, \tag{4.29}$$

where $r = q_1/q_0 < 1$ and $E_N(0) = 0$.

Proof. Given that the process $X^{(N)}(\cdot)$ is in state n at a point of the global clock, the probability with which it transits to the state $n+1$ at the next point of the global clock is denoted by $p(n \to n+1)$. This transition occurs when an agent with opinion $\{0\}$ updates its opinion to opinion $\{1\}$ by interacting with two agents having opinion $\{1\}$. The probability with which one among $N - n$ agents having opinion $\{0\}$ decides to update its opinion at a point of the global clock is $q_0(N - n)/N$. The probability with which the updating agent samples two agents with opinion $\{1\}$ is $(n/N)^2$. Hence, total probability of this transition is given by

$$p(n \to n+1) = \left(1 - \frac{n}{N}\right)\left(\frac{n}{N}\right)^2 q_0 \tag{4.30}$$

Similarly, the probabilities of other possible transitions are given by

$$p(n \to n-1) = \left(1 - \frac{n}{N}\right)^2\left(\frac{n}{N}\right) q_1 \tag{4.31}$$

$$p(n \to n) = 1 - p(n \to n+1) - p(n \to n). \tag{4.32}$$

Since $X^{(N)}(\cdot)$ is Markov, the exit probability $E_N(n)$ must satisfy the following recursive relationship:

$$E_N(n) = p(n \to n+1)E_N(n+1) + p(n \to n-1)E_N(n-1) + p(n \to n)E_N(n) \tag{4.33}$$

Using (4.30), (4.31), and (4.32) the above relationship can be simplified to

$$E_N(n+1) - E_N(n) = \frac{N-n}{n}\frac{q_1}{q_0}(E_N(n) - E_N(n-1)) \tag{4.34}$$

Putting $D_N(n) = E_N(n+1) - E_N(n)$ we find that (4.34) reduces to a first order recursion in $D_N(n)$ whose general solution is given by

$$D_N(n) = \left(\frac{B}{n!(N-n-1)!}\right)\left(\frac{q_1}{q_0}\right)^n \tag{4.35}$$

To compute the constant B we use the boundary conditions $E_N(0) = 0$ and $E_N(N) = 1$, which imply that $\sum_{n=0}^{N-1} D_N(n) = 1$. Hence, the constant B is given by

$$B = \frac{(N-1)!}{(1+r)^{N-1}}, \tag{4.36}$$

where $r = q_1/q_0$. Thus, we have

$$E_N(n) = \sum_{k=0}^{n-1} D_N(k) = \frac{(N-1)!}{(1+r)^{N-1}} \sum_{k=0}^{n-1} \frac{r^k}{k!(N-k-1)!}, \tag{4.37}$$

which simplifies to (4.29). □

We now proceed to characterize the mean consensus time of the network by analyzing the mean field limit of the empirical measure process $x^{(N)}(\cdot) = X^{(N)}(\cdot)/N$. It is easy to see that the rates of transition of the process $x(\cdot)$ from the state x to the states $x \pm 1/N$ are given by

$$r(x \to x + 1/N) = N(1-x)x^2 q_0 \tag{4.38}$$

$$r(x \to x - 1/N) = Nx(1-x)^2 q_1. \tag{4.39}$$

From the above transition rates, it easily follows that as $N \to \infty$ the generator of $x^{(N)}(\cdot)$ converges uniformly to the generator of the process $x(\cdot)$ satisfying the following differential equation:

$$\dot{x}(t) = x^2(t)(1-x(t))q_0 - x(t)(1-x(t))^2 q_1 = (q_0+q_1)x(t)(1-x(t))\left(x(t)-\kappa_q\right), \tag{4.40}$$

where $\kappa_q = q_1/(q_0+q_1)$. Therefore, Theorem C.2.2 and Theorem C.2.3 imply that, under the condition $x^{(N)}(0) \Rightarrow \alpha$, $x^{(N)}(\cdot) \Rightarrow x(\cdot)$ holds with $x(0) = \alpha$. In other words, the process $x^{(N)}(\cdot)$ weakly converges to the deterministic process $x(\cdot)$ as $N \to \infty$.

From (4.40), it follows that the process $x(\cdot)$ has three equilibrium points at 0, 1, and κ_q, respectively. We now characterize the stability of these equilibrium points in the sense of the following definition:

Definition 4.3.1. *An equilibrium point $x_e \in [0,1]$ of the process $x(\cdot)$ is called* stable *if there exists a non-empty set $S \subseteq [0,1]$ containing x_e but $S \neq \{x_e\}$ such that for all $x(0) \in S$ we have $x(t) \to x_e$ as $t \to \infty$. If no such sets exist, then x_e is called an* unstable *equilibrium point. The equilibrium point x_e is said to be globally stable if for all $x(0) \in [0,1]$ we have $x(t) \to x_e$ as $t \to \infty$.*

If for some $t \geq 0$ we have $1 \geq x(t) > \kappa_q$, then (4.40) implies $\dot{x}(t) \geq 0$. Hence, for $x(0) \in (\kappa_q, 1]$ we have $x(t) \to 1$ as $t \to \infty$. Similarly, for $x(0) \in [0, \kappa_q)$ we have $x(t) \to 0$ as $t \to \infty$. Therefore, $0, 1$ are the stable equilibrium points of the process $x(\cdot)$, and κ_q is an unstable equilibrium point.

If $x^{(N)}(0) = \alpha = \kappa_q + \epsilon$ ($\epsilon > 0$), then, for large N, with high probability the process $x^{(N)}(\cdot)$ reaches the state 1 in a finite time. This is because for large N the path of the process $x^{(N)}(\cdot)$ is close to that of $x(\cdot)$ with high probability (by the mean field convergence result) and we have already shown that for $x(0) = \alpha > \kappa_q$, $x(t) \to 1$ as $t \to \infty$. Therefore, the mean consensus time for large N and $\alpha = \kappa_q + \epsilon$ can be approximated as the time taken by the process $x(\cdot)$ to reach the state $1 - 1/N$ from the state $\alpha = \kappa_q + \epsilon$. We denote the approximate mean consensus time by $\bar{T}_N(\kappa_q + \epsilon)$. Now, solving (4.40) with the above limits we obtain

$$\bar{T}_N(\kappa_q + \epsilon) = \frac{1}{q_0 + q_1}\left[\frac{1}{\kappa_q(1-\kappa_q)}\ln(N(1-\kappa_q)-1) - \frac{1}{\kappa_q}\ln(N-1) - \frac{1}{\kappa_q(1-\kappa_q)}\ln\epsilon + \frac{1}{\kappa_q}\ln(\kappa_q+\epsilon) + \frac{1}{1-\kappa_q}\ln(1-\kappa_q-\epsilon)\right] \quad (4.41)$$

Similarly, for $\alpha = \kappa_q - \epsilon$, $\bar{T}_N(\alpha)$ can be approximated as the time taken by the process $x(\cdot)$ to reach the state $1/N$ from state α. Again by solving (4.40) with the above limits we have

$$\bar{T}_N(\kappa_q - \epsilon) = \frac{1}{q_0 + q_1}\left[\frac{1}{\kappa_q(1-\kappa_q)}\ln(N\kappa_q-1) - \frac{1}{1-\kappa_q}\ln(N-1) - \frac{1}{\kappa_q(1-\kappa_q)}\ln\epsilon + \frac{1}{\kappa_q}\ln(\kappa_q-\epsilon) + \frac{1}{1-\kappa_q}\ln(1-\kappa_q+\epsilon)\right] \quad (4.42)$$

From the above expressions, it is clear that the mean consensus time scales as $O(\ln N)$. We also observe that (4.41) can be obtained from (4.42) by replacing κ_q with $1-\kappa_q$. This symmetry is intuitive since interchanging q_0 and q_1 interchanges the statistical behaviors of agents having opinion $\{1\}$ and agents having opinion $\{0\}$.

Numerical Results: In Figure 4.5(a), we plot the exit probability as a function of the total number N of agents in the network. The parameters are chosen to be $q_0 = 1$, $q_1 = 0.6$. We observe that for $\alpha > \kappa_q$, the exit probability increases to 1 with the increase in N and for $\alpha < \kappa_q$, the exit probability decreases to zero with the increase in N. This implies that consensus is achieved on the preferred opinion (opinion $\{1\}$) with high probability only if the initial fraction of agents having the preferred opinion is more than the threshold given by κ_q. This is unlike the voter model with biased agents where the consensus is achieved on the preferred opinion always with a higher probability.

In Figure 4.5(b), we plot the exit probability as a function of the initial fraction α of agents having opinion $\{1\}$ for the same parameter setting. In this case, we observe a sharp increase in the exit probability in the range $\alpha \in (0.2, 0.6)$ which contains the unstable equilibrium equilibrium point $\kappa_q = 0.375$. This sharp change is due to the fact that for

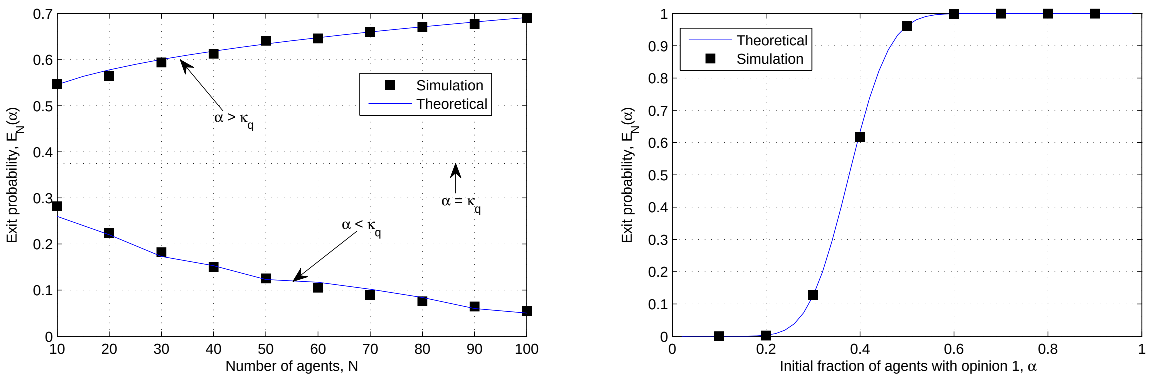

(a) Exit probability $E_N(\alpha)$ as a function of the number of agents N. Parameters: $q_0 = 1$, $q_1 = 0.6$.

(b) Exit probability $E_N(\alpha)$ as a function of the initial fraction α of agents with opinion $\{1\}$. Prameters: $q_0 = 1, q_1 = 0.6, N = 50$.

Figure 4.5: Exit probability under the majority rule with biased agents

$\alpha \in (\kappa_q, 1]$ (resp. $\alpha \in [0, \kappa_q)$) the process $X^{(N)}(\cdot)$ converges with high probability to the consensus state where all agents have opinion $\{1\}$ (resp. opinion $\{0\}$).

In Figure 4.6(a), we plot the mean consensus time as a function of N. Both simulation results and approximations obtained from (4.41) are plotted. We see that the mean consensus time slowly increases with increase in N, which is suggested by the $\ln(N)$ terms appearing in (4.41) and (4.42).

In Figure 4.6(b), we plot the mean consensus time as a function of the initial fraction of agents having opinion $\{1\}$. From the plot it is clear that the consensus time is maximum when the initial fraction of agents with opinion $\{1\}$ is near the unstable equilibrium point κ_q. This is because $\alpha = \kappa_q$ represents the critical point below which the process $x^{(N)}(\cdot)$ converges to the state 0 with high probability and above which it reaches the state 1 with high probability.

4.3.2 The majority rule with stubborn agents

We now consider the scenario where there are agents in the network who never update their opinions. We call these agents as the *stubborn* agents. The other agents, referred to as the *non-stubborn* agents, are assumed to update their opinions at all points of the Poisson processes associated with themselves. The updates occur according to the majority rule discussed in the beginning of this section. We denote by γ_i, $i \in \{0, 1\}$, the fraction of

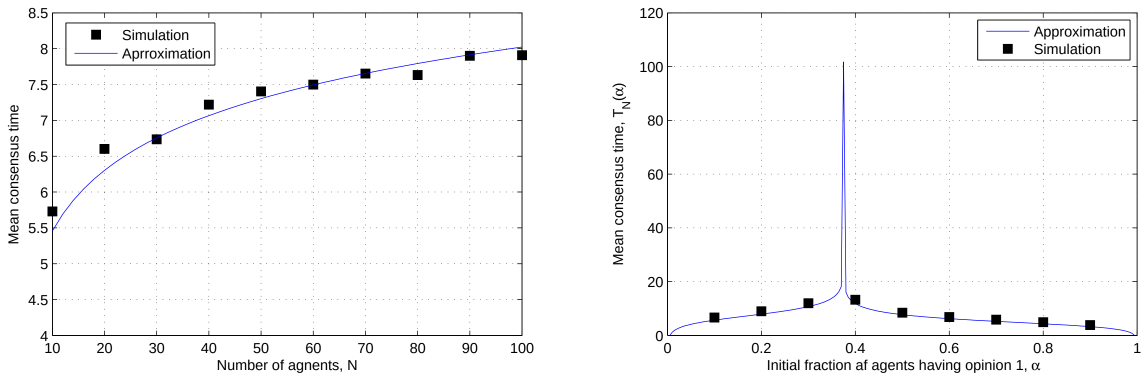

(a) Mean consensus time $T_N(\alpha)$ as a function of the number of agents N. Parameters: $q_0 = 1$, $q_1 = 0.6$, $\kappa_q = 0.375$, $\alpha = \kappa_q + \epsilon = 0.5$.

(b) Mean consensus time $T_N(\alpha)$ as a function of the initial fraction α of agents with opinion $\{1\}$. Prameters: $q_0 = 1, q_1 = 0.6, N = 200$.

Figure 4.6: Mean consensus time under the majority rule with biased agents

agents in network who are stubborn and have opinion i at all times. Thus, $(1 - \gamma_0 - \gamma_1)$ is fraction of non-stubborn agents in the network.

The presence of stubborn agents prevents the network from reaching a consensus state. This is because at all times there are at least $N\gamma_0$ stubborn agents having opinion $\{0\}$ and $N\gamma_1$ stubborn agents having opinion $\{1\}$. Furthermore, since each non-stubborn agent may interact with some stubborn agents at every update instant, it is always possible for the non-stubborn agent to change its opinion. Below we characterize the equilibrium fraction of non-stubborn agents having opinion $\{1\}$ in the network for large N.

Let $x^{(N)}(t)$ denote the fraction of non-stubborn agents having opinion $\{1\}$ at time $t \geq 0$. Clearly, $x^{(N)}(\cdot)$ is a Markov process with possible jumps at the points of a rate $N(1 - \gamma_0 - \gamma_1)$ Poisson process. The process $x^{(N)}(\cdot)$ jumps from the state x to the state $x + 1/N(1 - \gamma_0 - \gamma_1)$ when one of the non-stubborn agents having opinion $\{0\}$ becomes active (which happens with rate $N(1-\gamma_0-\gamma_1)(1-x)$) and samples two agents with opinion $\{1\}$. The probability of sampling an agent having opinion $\{1\}$ from the entire network is $(1 - \gamma_0 - \gamma_1)x + \gamma_1$. Hence, the total rate at which the process transits from state x to the state $x + 1/N(1 - \gamma_0 - \gamma_1)$ is given by

$$r\left(x \to x + \frac{1}{N(1 - \gamma_0 - \gamma_1)}\right) = N(1 - \gamma_0 - \gamma_1)(1 - x)[(1 - \gamma_0 - \gamma_1)x + \gamma_1]^2 \qquad (4.43)$$

Similarly, the rate of the other possible transition is given by

$$r\left(x \to x - \frac{1}{N(1-\gamma_0-\gamma_1)}\right) = N(1-\gamma_0-\gamma_1)x[(1-\gamma_0-\gamma_1)(1-x)+\gamma_0]^2 \quad (4.44)$$

It is easy to see from the above transition rates that the generator of the process $x^{(N)}(\cdot)$ converges uniformly to the generator of the process $x(\cdot)$ which satisfies the following differential equation

$$\dot{x}(t) = (1-x(t))[(1-\gamma_0-\gamma_1)x(t)+\gamma_1]^2 - x(t)[(1-\gamma_0-\gamma_1)(1-x(t))+\gamma_0]^2. \quad (4.45)$$

Hence, by applying Theorem C.2.2 and Theorem C.2.3, we have that if $x^{(N)}(0) \Rightarrow \alpha$ as $N \to \infty$, then $x^{(N)}(\cdot) \Rightarrow x(\cdot)$ as $N \to \infty$ with $x(0) = \alpha$. In other words, the process $x(\cdot)$ is the mean field limit of the sequence of processes $(x^{(N)}(\cdot))_N$. We now study the equilibrium distribution π_N of the process $x^{(N)}(\cdot)$ for large N via the equilibrium points of the mean field $x(\cdot)$.

From (4.45) we see that $\dot{x}(t)$ is a cubic polynomial in $x(t)$. Hence, the process $x(\cdot)$ can have at most three equilibrium points in $[0,1]$. We first characterize the stability of these equilibrium points according to Definition 4.3.1.

Proposition 4.3.2. *The process $x(\cdot)$ defined by (4.45) has at least one equilibrium point in $(0,1)$. Furthermore, the number of stable equilibrium points of $x(\cdot)$ in $(0,1)$ is either two or one. If there exists only one equilibrium point of $x(\cdot)$ in $(0,1)$, then the equilibrium point must be globally stable (attractive).*

Proof. Define

$$f(x) = (1-x)[(1-\gamma_0-\gamma_1)x+\gamma_1]^2 - x[(1-\gamma_0-\gamma_1)(1-x)+\gamma_0]^2 \quad (4.46)$$

Clearly, $f(0) = \gamma_1^2 > 0$ and $f(1) = -\gamma_0^2 < 0$. Hence, there exists at least one root of $f(x) = 0$ in $(0,1)$. This proves the existence of an equilibrium point of $x(\cdot)$ in $(0,1)$.

Since $f(x)$ is a cubic polynomial and $f(0)f(1) < 0$, either all three roots of $f(x) = 0$ lie in $(0,1)$ or exactly one root of $f(x) = 0$ lies in $(0,1)$. Let the three (possibly complex and non-distinct) roots of $f(x) = 0$ be denoted by r_1, r_2, r_3, respectively. By expanding (4.46) we see that the coefficient of the cubic term is $-2(1-\gamma_0-\gamma_1)^2$. Hence, $f(x)$ can be written as

$$f(x) = -2(1-\gamma_0-\gamma_1)^2(x-r_1)(x-r_2)(x-r_3) \quad (4.47)$$

We first consider the case when $0 < r_1, r_2, r_3 < 1$ and not all of them are equal.[4] Let us suppose, without loss of generality, that the roots are arranged in the increasing order, i.e., $0 < r_1 \leq r_2 < r_3 < 1$ or $0 < r_1 < r_2 \leq r_3 < 1$. From (4.47) and (4.45), it is clear that, if $x(t) > r_2$ and $x(t) > r_3$, then $\dot{x}(t) < 0$. Similarly, if $x(t) > r_2$ and $x(t) < r_3$, then $\dot{x}(t) > 0$. Hence, if $x(0) > r_2$ then $x(t) \to r_3$ as $t \to \infty$. Using similar arguments we have that for $x(0) < r_2$, $x(t) \to r_1$ as $t \to \infty$. Hence, r_1, r_3 are the stable equilibrium points of $x(\cdot)$. This proves that there exist at most two stable equilibrium points of the mean field $x(\cdot)$.

Now suppose that there exists only one equilibrium point of $x(\cdot)$ in $(0,1)$. This is possible either i) if there exists exactly one real root of $f(x) = 0$ in $(0,1)$, or ii) if all the roots of $f(x) = 0$ are equal and lie in $(0,1)$.[5] Let r_1 be a root of $f(x) = 0$ in $(0,1)$. Now by expanding $f(x)$ from (4.47), we see that the product of the roots must be $\gamma_1^2/2(1-\gamma_0-\gamma_1)^2 > 0$. This implies that the other roots, r_2 and r_3, must satisfy one of the following conditions:

- $r_2, r_3 > 1$.
- $r_2, r_3 < 0$.
- r_2, r_3 are complex conjugates.
- $r_2 = r_3 = r_1$.

In all the above cases, we have that $(x-r_2)(x-r_3) \geq 0$ for all $x \in [0,1]$ with equality if and only if $x = r_1 = r_2 = r_3$. Hence, from (4.47) and (4.45), it is easy to see that $\dot{x}(t) > 0$ when $0 \leq x(t) < r_1$ and $\dot{x}(t) < 0$ when $1 \geq x(t) > r_1$. This implies that $x(t) \to r_1$ for all $x(0) \in [0,1]$. In other words, r_1 is globally stable. □

In the next proposition, we provide the conditions on γ_0 and γ_1 for which there exist multiple stable equilibrium points of the mean field $x(\cdot)$.

Proposition 4.3.3. *There exist two distinct stable equilibrium points of the mean field $x(\cdot)$ in $(0,1)$ if and only if*

[4] To see that that such a situation is possible, consider the parameter setting $\gamma_0 = \gamma_1 = 0.2$. In this case, the three roots of $f(x) = 0$ are $x = 0.127322$, $x = 0.5$ and $x = 0.872678$.

[5] To see that this is possible consider the case where $\gamma_0 = 0.2$, $\gamma_1 = 0.3$.

1. $D(\gamma_0, \gamma_1) = (\gamma_0 - \gamma_1)^2 + 3(1 - 2\gamma_0 - 2\gamma_1) > 0$

2. $0 < z_1, z_2 < 1$*, where*

$$z_1 = \frac{(3 - \gamma_0 - 5\gamma_1) + \sqrt{D(\gamma_0, \gamma_1)}}{6(1 - \gamma_0 - \gamma_1)} \tag{4.48}$$

$$z_2 = \frac{(3 - \gamma_0 - 5\gamma_1) - \sqrt{D(\gamma_0, \gamma_1)}}{6(1 - \gamma_0 - \gamma_1)}. \tag{4.49}$$

3. $f(z_1)f(z_2) \leq 0$*, where* $f(x) = (1-x)[(1-\gamma_0-\gamma_1)x+\gamma_1]^2 - x[(1-\gamma_0-\gamma_1)(1-x)+\gamma_0]^2$.

If any one of the above conditions is not satisfied then $x(\cdot)$ *has a unique, globally stable equilibrium point in* $(0, 1)$*.*

Proof. From Proposition 4.3.2, we have seen that $x(\cdot)$ has two stable equilibrium points in $(0, 1)$ if and only if $f(x) = 0$ has three real roots in $(0, 1)$ among which at least two are distinct. This happens if and only if $f'(x) = 0$ has two distinct real roots z_1, z_2 in the interval $(0, 1)$ and $f(z_1)f(z_2) \leq 0$. Since $f'(x)$ is a quadratic polynomial in x, the above conditions are satisfied if and only if

1. The discriminant of $f'(x) = 0$ is positive. This corresponds to the first condition of the proposition.

2. The two roots z_1, z_2 of $f'(x) = 0$ must lie in $(0, 1)$. This corresponds to the second condition of the proposition.

3. $f(z_1)f(z_2) \leq 0$. This is the third condition of the proposition.

Clearly, if any one of the above conditions is not satisfied, then $x(\cdot)$ has a unique equilibrium point in $(0, 1)$. According to Proposition 4.3.2 this equilibrium point must be globally stable. □

In the above propositions, we have established that depending on the values of γ_0 and γ_1 there may exist of multiple stable equilibrium points of the mean field $x(\cdot)$. However, for every finite N, the process $x^{(N)}(\cdot)$ has a unique stationary distribution π_N (since it is irreducible on a finite state space). In the next result, we establish that any limit point of the sequence of stationary probability distributions $(\pi_N)_N$ is a convex combination of the Dirac measures concentrated on the equilibrium points of the mean field $x(\cdot)$ in $[0, 1]$.

Theorem 4.3.1. *Any limit point of the sequence of probability measures $(\pi_N)_N$ is a convex combination of the Dirac measures concentrated on the equilibrium points of $x(\cdot)$ in $[0,1]$. In particular, if there exists a unique equilibrium point r of $x(\cdot)$ in $[0,1]$ then $\pi_N \Rightarrow \delta_r$, where δ_r denotes the Dirac measure concentrated at the point r.*

Proof. We first note that since the sequence of probability measures $(\pi_N)_N$ is defined on the compact space $[0,1]$, it must be tight. Hence, Prokhorov's theorem implies that $(\pi_N)_N$ is relatively compact. Let π be any limit point of the sequence $(\pi_N)_N$. Then by the mean field convergence result we know that π must be an invariant distribution of the maps $\alpha \mapsto x(t,\alpha)$ for all $t \geq 0$, i.e.,

$$\int \varphi(x(t,\alpha))d\pi(\alpha) = \int \varphi(\alpha)d\pi(\alpha), \tag{4.50}$$

for all $t \geq 0$, and all continuous (and hence bounded) functions $\varphi : [0,1] \mapsto \mathbb{R}$. In the above, $x(t,\alpha)$ denotes the process $x(\cdot)$ started at $x(0) = \alpha$. Now, (4.50) implies

$$\int \varphi(\alpha)d\pi(\alpha) = \lim_{t\to\infty} \int \varphi(x(t,\alpha))d\pi(\alpha) \tag{4.51}$$

$$= \int \varphi\left(\lim_{t\to\infty} x(t,\alpha)\right) d\pi(\alpha) \tag{4.52}$$

The second equality follows from the first by the Dominated convergence theorem and the continuity of φ. Now, let r_1, r_2, and r_3 denote the three equilibrium points of the mean field $x(\cdot)$. Hence, by Proposition 4.3.2 we have that for each $\alpha \in [0,1]$, $\varphi(\lim_{t\to\infty} x(t,\alpha)) = \varphi(r_1)I_{N_{r_1}}(\alpha) + \varphi(r_2)I_{N_{r_2}}(\alpha) + \varphi(r_3)I_{N_{r_3}}(\alpha)$, where for $i = 1,2,3$, $N_{r_i} \in [0,1]$ denotes the set for which if $x(0) \in N_{r_i}$ then $x(t) \to r_i$ as $t \to \infty$, and I denotes the indicator function. Hence, by (4.52) we have that for all continuous functions $\varphi : [0,1] \mapsto \mathbb{R}$

$$\int \varphi(\alpha)d\pi(\alpha) = \varphi(r_1)\pi(N_{r_1}) + \varphi(r_2)\pi(N_{r_2}) + \varphi(r_3)\pi(N_{r_3}) \tag{4.53}$$

This proves that π must be of the form $\pi = c_1\delta_{r_1} + c_2\delta_{r_2} + c_3\delta_{r_3}$, where $c_1, c_2, c_3 \in [0,1]$ are such that $c_1 + c_2 + c_3 = 1$. This completes the proof. □

Thus, according to the above theorem, if there exists a unique equilibrium point of the process $x(\cdot)$ in [0,1], then the sequence of stationary distributions $(\pi_N)_N$ concentrates on that equilibrium point as $N \to \infty$. In other words, for large N, the fraction of non-stubborn

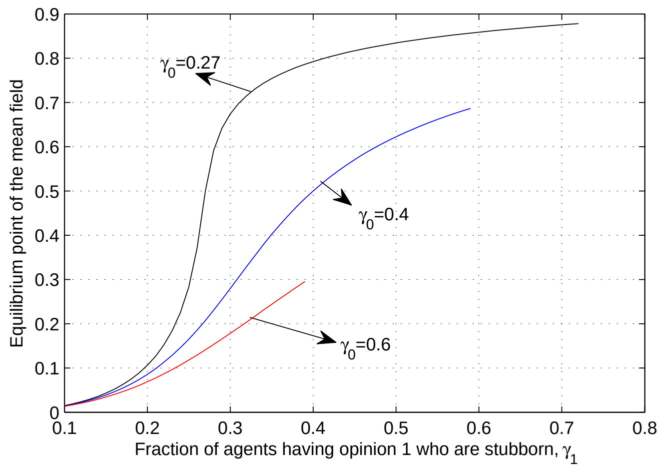

Figure 4.7: Majority rule with stubborn agents: Equilibrium point of $x(\cdot)$ as a function of γ_1 for different values of γ_0.

agents having opinion $\{1\}$ (at equilibrium) will approximately be equal to the unique equilibrium point of the mean field. We now numerically study the unique equilibrium point as a function of the fractions γ_0 and γ_1. In Figure 4.7, we plot the equilibrium point of $x(\cdot)$ as a function of the fraction γ_1 of agents having opinion $\{1\}$ who are stubborn keeping the fraction γ_0 of stubborn agents having opinion $\{0\}$ fixed. We choose the parameter values so that there exists a unique equilibrium point of $x(\cdot)$ in $[0, 1]$ (such parameter settings can be obtained using the conditions of Proposition 4.3.3). We see that as γ_1 is increased in the range $(0, 1 - \gamma_0 - \epsilon)$ $(\epsilon > 0)$, the equilibrium point shifts closer to unity. This is expected since increasing the fraction of stubborn agents with opinion $\{1\}$ increases the probability with which a non-stubborn agents samples an agent with opinion $\{1\}$ at an update instant.

If there exist multiple equilibrium points of the process $x(\cdot)$ then the convergence $x^{(N)}(\cdot) \Rightarrow x(\cdot)$ implies that at steady state the process $x^{(N)}(\cdot)$ spends intervals near the region corresponding to one of the stable equilibrium points of $x(\cdot)$. Then due to some rare events, it reaches, via the unstable equilibrium point, to a region corresponding to the other stable equilibrium point of $x(\cdot)$. This fluctuation repeats giving the process $x^{(N)}(\cdot)$ a unique stationary distribution.

The behavior discussed above is formally known as *metastability*. Although systems showing such metastable behavior are rare in practice, there are examples of such systems in the context of loss networks. See for example [55]. To demonstrate metastability, we

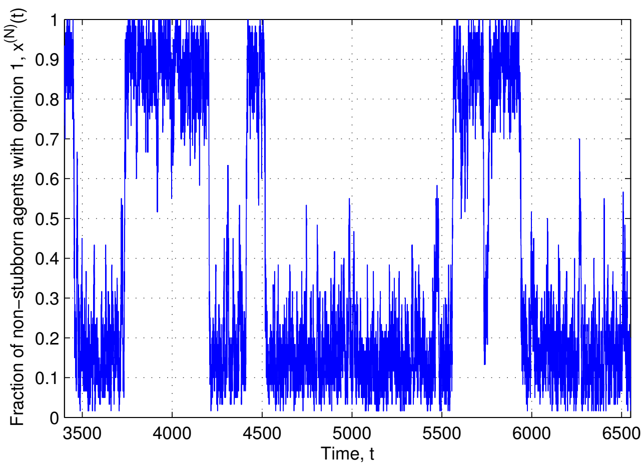

Figure 4.8: Majority rule with stubborn agents: Sample path of the process $x^{(N)}(\cdot)$ with $N = 100$, $\gamma_0 = \gamma_1 = 0.2$.

simulate a network with $N = 100$ agents and $\gamma_0 = \gamma_1 = 0.2$. For the above parameters, the mean field $x(\cdot)$ has two stable equilibrium points at 0.127322 and 0.872678. In Figure 4.8, we show the sample path of the process $x^{(N)}(\cdot)$. We see that at steady state the process switches back and forth between regions corresponding to the stable equilibrium points of $x(\cdot)$. This provides numerical evidence of the metastable behavior of the finite system.

4.3.3 Sampling more than two neighbors

In this section, we study the scenarios described in Subsections 4.3.1 and 4.3.2 for $K \geq 1$. In this case, the analysis is similar to those discussed above but involves more algebraic computations. For example, it can be easily verified that for general K, the mean field limit $x(\cdot)$ of the majority rule model with biased agents satisfies the following differential equation

$$\dot{x}(t) = \sum_{i=K+1}^{2K} x(t)(1-x(t))\binom{2K}{i}\left[x(t)^{i-1}(1-x(t))^{2K-i}q_0 - (1-x(t))^{i-1}x(t)^{2K-i}q_1\right]. \tag{4.54}$$

This implies that for any $K \geq 1$, two stable equilibrium points of the mean field are located at $0, 1$. The value of the unstable equilibrium however depends on the value of K.

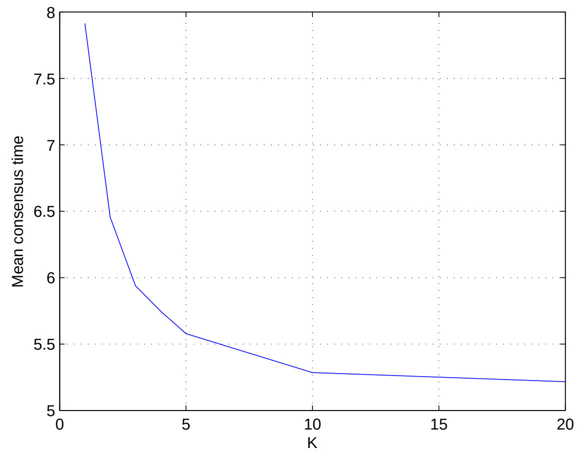

Figure 4.9: Majority rule with biased agents: Mean consensus time as a function of K. Parameters: $q_0 = 1, q_1 = 0.6, \alpha = 0.5, N = 50$

It turns out that increasing K (keeping q_0, q_1 unchanged) shifts the unstable equilibrium point towards one.

We first numerically study the majority rule model with biased agents for $K \geq 1$. It is expected that with the increase in K mean time to reach consensus will decrease. This is because for larger values of K the probability with which the majority opinion of a selected group of $2K + 1$ agents is the same as the majority opinion of whole network is higher. Therefore, when the state of the system is above (or below) the unstable equilibrium point of the mean field, the network experiences a stronger 'pull' towards the majority resulting in a smaller mean consensus time. In Figure 4.9 we plot the mean consensus time as a function of K for $q_0 = 1, q_1 = 0.6, \alpha = 0.5, N = 50$. As expected we observe that the mean consensus time decreases with the increase in K. However, since the dynamics of the majority rule follows that of a diffusion process, the asymptotic order ($\ln N$) of the consensus time remains unchanged.

Next, we study the majority rule model with stubborn agents for $K \geq 1$. We consider the case $\gamma_1 > \gamma_0$. In this case, with the increase in K, the probability with which the majority opinion in a group of $2K + 1$ agents is $\{1\}$, increases. Therefore, increase in K should result in a shift of the equilibrium fraction of non-stubborn agents with opinion $\{1\}$ closer to unity. In Figure 4.10 we plot the equilibrium fraction of non-stubborn agents having opinion $\{1\}$ as a function K for $\gamma_0 = 0.2, \gamma_1 = 0.4, N = 50$. As expected we observe that the equilibrium fraction of agents with opinion $\{1\}$ increases to unity as K increases.

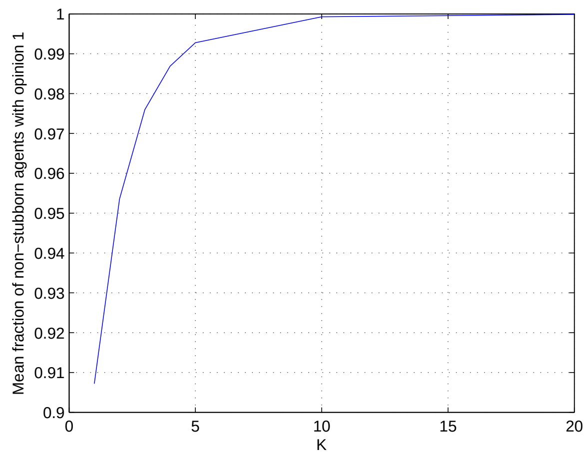

Figure 4.10: Majority rule with stubborn agents: Mean consensus time as a function of K. Parameters: $\gamma_0 = 0.2, \gamma_1 = 0.4, \alpha = 0.5, N = 50$

4.4 Conclusion

In this chapter, we analyzed models of opinion dynamics based on the voter rule and the majority rule. Under each rule, we considered scenarios i) where the agents are biased towards a specific opinion and ii) where the agents have different propensities to change their past opinions. We observed that for the voter model, the presence of biased agents, reduces the mean consensus time exponentially in comparison to the voter model with unbiased agents. For the majority rule model with biased agents, we saw that the network reaches the consensus state with all agents adopting the preferred opinion only if the initial fraction of agents having the preferred opinion is more than a certain threshold value. The threshold was computed in the large system limit using the mean field analysis. For the voter model with agents having different propensities to change their opinions, we obtained a closed form expression for the probability with which consensus is reached on a specific opinion. It was observed that this probability does not depend on the number of agents in the network. Finally, we have seen that for the majority rule model with stubborn agents the network exhibits metastability, where it fluctuates between multiple stable configuration, spending long intervals in each configuration.

Chapter 5

Conclusion

In this dissertation, we analyzed large systems of weakly interacting heterogeneous entities using mean field techniques. We focused on two instances of such systems arising in two different contexts. First, we focused on heterogeneous multi-server systems which serve as models of large web server farms and cloud data centers. Job assignment schemes, in which the incoming job requests are assigned to the servers by comparing the states of some randomly sampled servers at each job arrival instant, were analyzed. In the second part of the dissertation, we considered the effect of random interactions between agents in a social network on the diffusion of opinions in the network.

In the first part, both processor sharing systems and loss systems were studied. In Chapter 2, we considered randomized dynamic job assignment schemes for a system of parallel processor sharing servers having different capacities. It was shown that uniform sampling of the servers at each job arrival instant may lead to a reduction in the stability region of the system. A hybrid scheme, that combines biased sampling of different server types with uniform sampling of servers within the same type, was shown to achieve the maximal stability region with the knowledge of the system parameters. To avoid requiring the knowledge of the system parameters, a type-based scheme was proposed and shown to achieve the maximal stability region under all values of the system parameters. The performance of the type-based scheme was analyzed using mean field techniques. Numerical studies were conducted to show that the hybrid scheme and the type based scheme significantly outperform randomized static schemes where job assignment decisions are made independently of the server states.

In Chapter 3, randomized dynamic job assignment schemes were studied with the aim of reducing the average blocking probability of jobs arriving at a system of parallel loss servers. A job assignment scheme, where every arriving job is assigned to the server with the maximum available resource among a set of randomly sampled servers, was shown to have near optimal performance in terms of reducing the average blocking probability of jobs in the system. The performance of the scheme was analyzed using mean field techniques. A recursive formula to efficiently compute the average blocking probability of jobs was obtained for the multi-class case.

In Chapter 4, diffusion of binary opinions in a large social network was studied under the voter model and the majority rule model. Two scenarios were considered. In the first scenario, the agents were assumed to be biased towards a specific opinion. In this scenario, the voter model was shown to converge to a consensus state in a time that is logarithmic in the total number of agents. For the majority rule model with biased agents it was shown that all the agents eventually adopt the 'preferred' opinion with high probability only if the initial fraction of agents having the preferred opinion is above a certain threshold. A second scenario in which different agents have different propensities to change their opinions was also analyzed under the voter model and the majority rule model. For the voter model, the exit probability was found to be independent of the network size. For the majority rule model, we observed that the presence of stubborn agents results in a metastable behavior of the system under certain parameter settings.

5.1 Future Extensions

Following the results presented in this dissertation, there are several interesting avenues of future research. Some of the them are discussed below.

We have noted in Chapter 2 and Chapter 3 that in the large system limit the systems of processor sharing servers and loss servers operating under the randomized dynamic job assignment schemes become asymptotically insensitive to job length distributions. This is due to the following facts: i) In the limiting system the servers become independent of each other; ii) Each server has a symmetric service discipline (either processor sharing or loss servers). We have established asymptotic independence of servers only for the exponential job length distribution. Proving asymptotic independence for more general job length distributions remains as an open problem. The proof will require construction of Markov processes on continuous state spaces and analyzing their corresponding mean field limits.

In the model considered in Chapter 3, we assumed that each incoming job is assigned to the server having the maximum vacancy among a randomly sampled set of d servers. An interesting variation of the scheme would be a scheme where an incoming job is assigned to the server (among the set of sampled servers) whose vacancy 'matches' the 'best' with the resource requirement of the job. This will ensure that a job with a 'small' resource requirement will not occupy a server with a much 'larger' vacancy where a future potential arrival having a larger resource requirement can be assigned. Such a criterion of server selection, based on 'best fit', is expected to yield lower blocking probability of jobs than the scheme considered in Chapter 3. However, a more precise definition of the scheme is required and a detailed analysis need to be carried out to verify this.

We assumed in Chapter 3 that the resource requirements of the incoming jobs are one-dimensional in the sense that different classes of jobs require different units of the same resource. A model where resource requirements are multi-dimensional would be a more accurate model for cloud computing systems, where each incoming job requires specific amounts of different resources (e.g. memory, CPU, storage).

In the majority rule model considered in Chapter 4, we assumed that each agent samples a fixed number of agents in the network at every update instant. A generalization of this model would one in which the agents sample random number of agents with a fixed mean. In such a scenario, it would be interesting to study how the mean consensus time depends on the distribution of the number of agents sampled during an update instant.

Another scenario would be one in which the agents have fixed neighborhoods at all times. In this case, analyzing the voter model and the majority rule model assuming the presence of biased and stubborn agents is much more challenging. Mean field techniques cannot be applied in such a scenario. A generalization of coalescing random walk techniques may provide bounds on the mean consensus time in such cases.

Appendices

Appendix A

Proof of compactness of $\bar{\mathcal{U}}^M$

We prove that the space $\bar{\mathcal{U}}^M$, defined as the M fold product of the space

$$\bar{\mathcal{U}} = \{(g_n, n \in \mathbb{Z}_+) : g_0 = 1, g_n \geq g_{n+1} \geq 0 \text{ for all } n \in \mathbb{Z}_+\}, \tag{A.1}$$

is compact under the metric ω, defined as

$$\omega(\mathbf{u}, \mathbf{v}) = \sup_{1 \leq j \leq M} \sup_{k \in \mathbb{Z}_+} \left| \frac{u_{k,j} - v_{k,j}}{k+1} \right|, \text{ for all } \mathbf{u}, \mathbf{v} \in \bar{\mathcal{U}}^M. \tag{A.2}$$

It is sufficient to show that the space $\bar{\mathcal{U}}$ (defined in (A.1)) is compact under the metric $\tilde{\omega}$ defined as

$$\tilde{\omega}(u, v) = \sup_{k \in \mathbb{Z}_+} \left| \frac{u(k) - v(k)}{k+1} \right|, \tag{A.3}$$

where $u = (u(k), k \in \mathbb{Z}_+), v = (v(k), k \in \mathbb{Z}_+) \in \bar{\mathcal{U}}$.

In order to prove the compactness of $\bar{\mathcal{U}}$ under the metric $\tilde{\omega}$, we prove that any sequence $(u_n)_n \subset \bar{\mathcal{U}}$ has a sub-sequence $(u_{n_m})_m \subset (u_n)_n$ convergent to a limit $u \in \bar{\mathcal{U}}$ under the metric $\tilde{\omega}$. We first note that for each $k \in \mathbb{Z}_+$, the sequence $(u_n(k))_n$ is a bounded sequence in $[0, 1]$ and therefore has a convergent sub-sequence. Hence, by the process of diagonalization, we can find a sub-sequence $(u_{n_m})_m$ of $(u_n)_n$ such that for each $k \in \mathbb{Z}_+$, $u_{n_m}(k) \to u(k)$ as $m \to \infty$ for some $u(k) \in [0, 1]$. Clearly, $u = (u(k), k \in \mathbb{Z}_+) \in \bar{\mathcal{U}}$. We now show that $\tilde{\omega}(u_{n_m}, u) \to 0$ as $m \to \infty$.

Since for each $k \in \mathbb{Z}_+$ we have $u_{n_m}(k) \to u(k)$ as $m \to \infty$, we can choose m sufficiently large such that $\left|\frac{u_{n_m}(k)-u(k)}{k+1}\right| < \frac{1}{l+1}$ for $0 \le k \le l$. For $k > l$, we have $\left|\frac{u_{n_m}(k)-u(k)}{k+1}\right| < \frac{1}{l+1}$ since $|u_{n_m}(k) - u(k)| \le 1$ and $k > l$. Thus, we have that for sufficiently large m

$$\tilde{\omega}(u_{n_m}, u) = \sup_{k \in \mathbb{Z}_+} \left|\frac{u(k) - v(k)}{k+1}\right| < \frac{1}{l+1}. \tag{A.4}$$

Since the above holds for any $l \in \mathbb{N}$, the required convergence follows.

Appendix B

Proof of Theorem 2.6.3

We provide a proof for $M = 2$. The proof can be generalized to any $M \geq 2$.

The main idea is to construct sequences $\{P_{k,j}, k \in \mathbb{Z}_+\}$ for $j = 1, 2$ such that they satisfy the following three properties

P.1 Equation (2.85) for $j = 1, 2$.

P.2 $P_{0,j} = 1 \geq P_{k,j} \geq 0$ for all $k \in \mathbb{Z}_+$, $j = 1, 2$.

P.3 $P_{k,j} \to 0$ as $k \to \infty$ for $j = 1, 2$.

The first property ensures that the constructed $\mathbf{P} = \{P_{k,j}, k \in \mathbb{Z}_+, j \in \{1, 2\}\}$ satisfies $\mathbf{l}(\mathbf{P}) = \mathbf{0}$. The second property (in conjunction with the first property) implies that $P_{k,j} \geq P_{k+1,j} \geq 0$ for all $k \in \mathbb{Z}_+$, $j = 1, 2$. Hence, $\mathbf{P} \in \bar{\mathcal{U}}^2$. Finally, according to Proposition 2.6.2 the third property (along with the first two properties) guarantees that $\mathbf{P}$ must lie in the space $\mathcal{U}^2$. Hence the proof is complete if we can construct a $\mathbf{P}$ satisfying all the three properties.

We now construct the sequences $\{P_{l,1}(\alpha), l \in \mathbb{Z}_+\}$ and $\{P_{l,2}(\alpha), l \in \mathbb{Z}_+\}$ as functions of the real variable α as follows:

$$P_{0,1}(\alpha) = 1. \tag{B.1}$$
$$P_{0,2}(\alpha) = 1. \tag{B.2}$$
$$P_{1,1}(\alpha) = \alpha. \tag{B.3}$$
$$P_{1,2}(\alpha) = \Delta_2 \left(1 - \frac{\alpha}{\Delta_1}\right). \tag{B.4}$$
$$P_{l+2,1}(\alpha) = P_{l+1,1}(\alpha) - \Delta_1 \left((P_{l,1}(\alpha))^{d_1} - (P_{l+1,1}(\alpha))^{d_1} \right) \times (P_{l+1,2}(\alpha))^{d_2}, l \geq 0 \tag{B.5}$$
$$P_{l+2,2}(\alpha) = P_{l+1,2}(\alpha) - \Delta_2 \left((P_{l,2}(\alpha))^{d_2} - (P_{l+1,2}(\alpha))^{d_2} \right) \times (P_{l,1}(\alpha))^{d_1}, l \geq 0 \tag{B.6}$$

Combining the above relations we obtain

$$\sum_{j=1}^{2} \frac{P_{l+1,j}(\alpha)}{\Delta_j} = \prod_{j=1}^{2} (P_{l,j}(\alpha))^{d_j}, \text{ for } l \geq 0 \tag{B.7}$$

Note that the sequences $\{P_{l,1}(\alpha), l \in \mathbb{Z}_+\}$ and $\{P_{l,2}(\alpha), l \in \mathbb{Z}_+\}$ are constructed such that they satisfy property (P.1). Hence, the the proof will be complete if for some α the properties (P.2) and (P.3) are satisfied. We first proceed to find the range of α for which the sequences $\{P_{l,1}(\alpha), l \in \mathbb{Z}_+\}$ and $\{P_{l,2}(\alpha), l \in \mathbb{Z}_+\}$ are both positive sequences of real numbers in $[0, 1]$. This will ensure that (P.2) is satisfied.

We first observe from (B.3) that to have $1 \geq P_{1,1}(\alpha) \geq 0$ we must choose α in the range $[0, 1]$. Also, from (B.4) we observe that to have $1 \geq P_{1,2}(\alpha) \geq 0$ we must choose $\alpha \in [\Delta_1(1 - 1/\Delta_2), \Delta_1]$. Combining the above two ranges we have the following effective range of α for which $1 \geq P_{1,1}(\alpha), P_{2,1}(\alpha) \geq 0$

$$\alpha \in \left[\max\left(0, \Delta_1 \left(1 - \frac{1}{\Delta_2}\right)\right), \min(1, \Delta_1)\right]. \tag{B.8}$$

We note that the above range is always non-empty due to the stability condition $\lambda \in \Lambda$ which implies that $1/\Delta_1 + 1/\Delta_2 > 1$.

The above range can be further refined as follows. From (B.4), (B.6), and (B.3) we see that at $\alpha = 0$, $P_{1,1}(0) = 0$, $P_{1,2}(0) = \Delta_2 > 0$, $P_{l,2} = \Delta_2^{d_2+1} > 0$ for all $l \geq 2$.

Hence, at $\alpha = 0$, $P_{l,2}(0) > 0$ for all $l \geq 0$. From the same relations we also observe that $P_{l,2}\left(\Delta_1\left(1 - \frac{1}{\Delta 2}\right)\right) = 1 > 0$ for all $l \geq 0$. Combining the two we have that for all $l \geq 0$

$$P_{l,2}\left(\max\left(0, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right)\right) > 0 \tag{B.9}$$

Now, putting $l = 0$, $\alpha = \Delta_1$ in (B.6) we observe that $P_{2,2}(\Delta_1) < 0$. Hence, from (B.9) it follows that there exists at least one root of $P_{2,2}(\alpha)$ in the range $\left[\max\left(0, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \Delta_1\right]$. Let $r_{2,2}$ denote the minimum of these roots. Then in the range

$$\alpha \in \left[\max\left(0, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{2,2}\right)\right], \tag{B.10}$$

we have $1 \geq P_{1,1}(\alpha) \geq 0, 1 \geq P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \geq 0$. Repeating the same argument again for $P_{3,2}(\alpha)$ we find that $1 \geq P_{1,1}(\alpha) \geq 0$ and $1 \geq P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \geq P_{3,2}(\alpha) \geq 0$ in the range

$$\alpha \in \left[\max\left(0, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3,2}\right)\right], \tag{B.11}$$

where $r_{3,2}$ denotes the minimum root of $P_{3,2}(\alpha)$ in the range defined in (B.10).

We now refine the left limit of the above range as follows. From (B.5) we have $P_{2,1}(0) = -\Delta_1\Delta_2^{d_2} < 0$. Also, from definition of $r_{3,2}$ we know that $P_{3,2}(r_{3,2}) = 0$. Now, by putting $\alpha = r_{3,2}$ and $l = 1$ in (B.6) we obtain

$$\begin{aligned} P_{2,2}(r_{3,2}) &= \Delta_2\left[\left(P_{1,2}(r_{3,2})\right)^{d_2} - \left(P_{2,2}(r_{3,2})\right)^{d_2}\right]\left(r_{3,2}\right)^{d_1} \\ &\leq \Delta_2\left(P_{1,2}(r_{3,2})\right)^{d_2}\left(r_{3,2}\right)^{d_1} \text{ (since } P_{2,2}(r_{3,2}) \geq 0) \end{aligned}$$

Again, by putting $l = 2$ and $\alpha = r_{3,2}$ in (B.7) and using the above we obtain $P_{2,1}(r_{3,2}) \geq 0$. Therefore, there exists at least one root of $P_{2,1}(\alpha)$ in the interval $(\,0, r_{3,2}\,]$. Denote the maximum of all such roots to be $r_{2,1}$. Hence, in the range

$$\alpha \in \left[\max\left(r_{2,1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3,2}\right)\right], \tag{B.12}$$

we have $1 \geq P_{1,1}(\alpha) \geq P_{2,1}(\alpha) \geq 0$ along with $1 \geq P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \geq P_{3,2}(\alpha) \geq 0$. Again from (B.5) we observe that $P_{3,1}(r_{2,1}) < 0$. Further, putting $l = 3$ and $\alpha = r_{3,2}$ in (B.7) we obtain $P_{3,1}(r_{3,2}) \geq 0$. Thus, there must be at least one root of $P_{3,1}(\alpha)$ in the range $(\, r_{2,1}, r_{3,2}\,]$. Let $r_{3,1}$ denote the maximum root in the interval. Hence, in the interval

$$\alpha \in \left[\max\left(r_{3,1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3,2}\right)\ \right], \tag{B.13}$$

we have $1 \geq P_{1,1}(\alpha) \geq P_{2,1}(\alpha) \geq P_{3,1}(\alpha) \geq 0$ along with $1 \geq P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \geq P_{3,2}(\alpha) \geq 0$. Similarly, from (B.5) we have $P_{4,1}(r_{3,1}) < 0$ and from (B.6) we have $P_{4,1}(r_{3,2}) \geq 0$. Thus, there must be at least one root of $P_{4,1}(\alpha)$ in the range $(\, r_{3,1}, r_{3,2}\,]$. Denote the maximum of all such roots by $r_{4,1}$. Hence, in the interval

$$\alpha \in \left[\max\left(r_{4,1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3,2}\right)\ \right], \tag{B.14}$$

we have $1 \geq P_{1,1}(\alpha) \geq P_{2,1}(\alpha) \geq P_{3,1}(\alpha) \geq P_{4,1}(\alpha) \geq 0$ and $1 \geq P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \geq P_{3,2}(\alpha) \geq 0$.

Using the same line of arguments as above, the following inductive hypothesis can be proved: If $P_{1,1}(\alpha) \geq P_{2,1}(\alpha) \ldots \geq P_{4+3k,1}(\alpha) \geq 0$ and $P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \ldots \geq P_{3+3k,1}(\alpha) \geq 0$ for some $k \geq 0$ and

$$\alpha \in \left[\max\left(r_{4+3k,1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3+3k,2}\right)\ \right], \tag{B.15}$$

then $P_{1,1}(\alpha) \geq P_{2,1}(\alpha) \ldots \geq P_{4+3(k+1),1}(\alpha) \geq 0$ and $P_{1,2}(\alpha) \geq P_{2,2}(\alpha) \ldots \geq P_{3+3(k+1),1}(\alpha) \geq 0$ hold in the range

$$\alpha \in \left[\max\left(r_{4+3(k+1),1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3+3(k+1),2}\right)\ \right], \tag{B.16}$$

and the interval in (B.16) is included in the interval in (B.15).

The of compact intervals

$$\left[\max\left(r_{4+3k,1}, \Delta_1\left(1 - \frac{1}{\Delta_2}\right)\right), \min\left(1, r_{3+3k,2}\right)\ \right], k \geq 0 \tag{B.17}$$

eventually become strict subsets of the interval $[0, 1]$. This can be justified as follows. We note that $P_{l,1}(1) = 1$ for all $l \in \mathbb{Z}_+$. Hence, from (B.4) we have $P_{1,2}(1) = \Delta_2\left(1 - \frac{1}{\Delta_1}\right)$ and from (B.6) we have

$$P_{l+2,2}(1) = P_{l+1,2}(1) - \Delta_2 \left((P_{l,2}(1))^{d_2} - (P_{l+1,2}(1))^{d_2} \right) \text{ for } l \geq 0 \tag{B.18}$$

Notice that the stability condition $\lambda \in \Lambda$ can be equivalently expressed as

$$\frac{1}{\Delta_1} + \frac{1}{\Delta_2} > 1, \tag{B.19}$$

which implies that $P_{1,2}(1) < 1$. We claim that there exists some $l \geq 1$ such that $P_{l,2}(1) < 0$. Let us assume this is not true. Therefore, $P_{l,2}(1) \geq 0$ for all $l \geq 0$. By (B.18), this implies that $\{P_{l,2}(1), l \geq 0\}$ is a non-increasing sequence of numbers in $[0, 1)$. Hence by monotone convergence theorem $\lim_{l\to\infty} P_{l,2}(1)$ exists. Let this limit be denoted by β, where $0 \leq \beta < 1$. Thus, adding (B.18) for $l \geq 0$ and using $\lim_{l\to\infty} P_{l,2}(1) = \beta$ we obtain

$$\begin{aligned}\left(1 - \frac{1}{\Delta_1}\right) &= \frac{\beta}{\Delta_2} + 1 - \beta^{d_2} \\ &> \beta \left(1 - \frac{1}{\Delta_1}\right) + 1 - \beta^{d_2}.\end{aligned}$$

Hence, $\left(1 - \frac{1}{\Delta_1}\right) > \frac{1-\beta^{d_2}}{1-\beta} \geq 1$. This is a contradiction since $\Delta_1 > 0$. Hence, there exists $l \geq 1$ such that $P_{l,2}(1) < 0$. This implies that for some $k \geq 0$, $r_{3k+3,2} < 1$. Similarly, by observing that $P_{l,2}\left(\Delta_1\left(1 - \frac{1}{\Delta 2}\right)\right) = 1$ for all $l \geq 0$, it can be shown that there exists $l \geq 1$ such that $P_{l,1}\left(\Delta_1\left(1 - \frac{1}{\Delta 2}\right)\right) < 0$. This implies there exists $k \geq 0$ for which $r_{4+3k,1} > \Delta_1(1 - 1/\Delta_2)$.

Further, the intersection of all such compact intervals must be non-empty due to the Cantor's intersection theorem. Hence, we have shown that there exists $\alpha \in (0, 1)$ such that the sequences $\{P_{l,1}(\alpha), l \in \mathbb{Z}_+\}$ and $\{P_{l,2}(\alpha), l \in \mathbb{Z}_+\}$ are both positive non-increasing sequences of real numbers in $[0, 1]$.

We now show that the above sequences satisfy property (P.3). Let $\lim_{l\to\infty} P_{l,1}(\alpha) = \xi_1 \geq 0$ and $\lim_{l\to\infty} P_{l,2}(\alpha) = \xi_2 \geq 0$, where $\alpha \in (0, 1)$ is chosen such that both sequences become positive and non-increasing. Now, taking limit of (B.7) as $l \to \infty$ we have

$$\sum_{j=1}^{2} \frac{\xi_j}{\Delta_j} = \prod_{j=1}^{2} (\xi_j)^{d_j}. \tag{B.20}$$

Now using the stability criterion and the fact that $0 \leq \xi_1, \xi_2 \leq 1$ we have

$$\frac{1}{\Delta_1} + \frac{1}{\Delta_2} > 1$$
$$\Rightarrow \frac{\xi_2}{\Delta_1} + \frac{\xi_2}{\Delta_2} \geq \xi_2 \geq \xi_2^{d_2}$$

with equality holding if and only if $\xi_2 = 0$. Further, we have

$$\frac{1}{\Delta_1} + \frac{\xi_2}{\Delta_2} \geq \frac{\xi_2}{\Delta_1} + \frac{\xi_2}{\Delta_2} \geq \xi_2^{d_2}$$

Hence, by multiplying both sides with ξ_1 we have

$$\frac{\xi_1}{\Delta_1} + \frac{\xi_1 \xi_2}{\Delta_2} \geq \xi_1 \xi_2^{d_2} \geq \xi_1^{d_1} \xi_2^{d_2},$$

with equality if and only if $\xi_1 = \xi_2 = 0$. Again, since $\xi_1 \leq 1$ we have

$$\frac{\xi_1}{\Delta_1} + \frac{\xi_2}{\Delta_2} \geq \frac{\xi_1}{\Delta_1} + \frac{\xi_1 \xi_2}{\Delta_2} \geq \xi_1 \xi_2^{d_2} \geq \xi_1^{d_1} \xi_2^{d_2},$$

Hence, we have shown

$$\frac{\xi_1}{\Delta_1} + \frac{\xi_2}{\Delta_2} \geq \xi_1^{d_1} \xi_2^{d_2} \tag{B.21}$$

with equality holding if and only if $\xi_1 = \xi_2 = 0$. Hence, for (B.20) to hold we must have $\xi_1 = \xi_2 = 0$. This proves (P.3) and thus completes the proof.

Appendix C

Mathematical Preliminaries

In the appendix, we review some of the key concepts and results related to weak convergence of probability measures and Markov processes. The results discussed here are used throughout this dissertation. For a more detailed discussion on the topics covered in this section the reader is referred to Chapters 1, 3, and 4 of [28].

Throughout we shall use (E, r) to denote a metric space; $\mathcal{B}(E)$ to denote the σ-algebra of Borel subsets of E; $\mathcal{P}(E)$ to denote the set of (Borel) probability measures on $(E, \mathcal{B}(E))$; $\bar{C}(E)$ to denote the space of real valued bounded continuous functions defined on E with the norm $\|f\|_\infty = \sup_{x \in E} |f(x)|$ for $f \in \bar{C}(E)$.

C.1 Weak convergence of probability measures

We begin by defining weak convergence of a sequence of probability measures in $\mathcal{P}(E)$.

Definition C.1.1. *A sequence of probability measures $\{P_n, n \geq 1\} \in \mathcal{P}(E)$ is said to converge weakly to $P \in \mathcal{P}(E)$ (written as $P_n \Rightarrow P$) if for all $f \in \bar{C}(E)$*

$$\lim_{n \to \infty} \int f dP_n = \int f dP. \tag{C.1}$$

Next, we define a metric d_W on $\mathcal{P}(E)$ with respect to which weak convergence of probability measures can be studied.

Definition C.1.2. *For $P, Q \in \mathcal{P}(E)$, the Prohorov metric $d_W(P, Q)$ is defined as*

$$d_W(P,Q) = \inf \{\epsilon > 0 : P(F) \leq Q(F^\epsilon) + \epsilon \text{ for all } F \in \mathcal{C}\}, \tag{C.2}$$

where $\mathcal{C}$ is the collection of closed subsets of E and for each $F \in \mathcal{C}$ we denote by F^ϵ the set $F^\epsilon = \{x \in E : \inf_{y \in F} d(x, y) < \epsilon\}$.

It can be shown that d_W is indeed a metric on $\mathcal{P}(E)$ (Lemma 1.1 of Chapter 3 of [28]). The following theorem establishes that a sequence of Borel probability measures on a separable metric space weakly converges to a limit if and only if it converges to the limit in the Prohorov's metric.

Theorem C.1.1 ([28], p. 108, Theorem 3.1)**.** *Let $\{P_n, n \geq 1\}$ be a sequence in $\mathcal{P}(E)$ and let $P \in \mathcal{P}(E)$. If (E, r) is separable then $P_n \Rightarrow P$ if and only if $lim_{n\to\infty} d_W(P_n, P) = 0$.*

In many cases, a direct proof of weak convergence of a sequence $\{P_n, n \geq 1\} \in \mathcal{P}(E)$ to a limiting probability measure $P \in \mathcal{P}(E)$ becomes difficult. However, if it is known that $\{P_n, n \geq 1\}$ is *relatively compact* (i.e., the closure of $\{P_n, n \geq 1\}$ in $\mathcal{P}(E)$ is compact), then to establish $P_n \Rightarrow P$ one just has to show that all convergent subsequences of $\{P_n, n \geq 1\}$ converge to the same limit P. The following theorem, due to Prohorov, provides a necessary and sufficient condition for a sequence of Borel probability measures to be relatively compact. Before stating the theorem we define the following notion of tightness of Borel probability measures.

Definition C.1.3. *A Borel probability measure $P \in \mathcal{P}(E)$ is said to be tight if for every $\epsilon > 0$ there exists a compact set $K \subset E$ such that $P(K) \geq 1 - \epsilon$. A family of Borel probability measures $\mathcal{M} \subset \mathcal{P}(E)$ is said to be tight if for every $\epsilon > 0$ there exists a compact set $K \subset E$ such that $\inf_{P \in \mathcal{M}} P(K) \geq 1 - \epsilon$.*

The following theorem establishes the equivalence between relative compactness and tightness of a set of Borel probability measures defined on a complete and separable metric space.

Theorem C.1.2 ([28], p. 104, Theorem 2.2)**.** *Let (E, r) be complete and separable and let $\mathcal{M} \subset \mathcal{P}(E)$. Then the following statements are equivalent*

(i) $\mathcal{M}$ is tight.

(ii) For each $\epsilon > 0$, there exists a compact set $K \subset E$ such that

$$\inf_{P \in \mathcal{M}} P(K^\epsilon) \geq 1 - \epsilon, \tag{C.3}$$

where K^ϵ is as defined in Definition C.1.2.

(iii) $\mathcal{M}$ is relatively compact.

The above theorem implies that a sequence of Borel probability measures on a complete and separable metric space is relatively compact if and only if it is tight. Note that if (E, r) is a compact metric space (hence complete and separable) then by definition every sequence of Borel probability measures is tight (to see this put $K = E$ in Definition C.1.3). Therefore, Prohorov's theorem implies that every sequence of Borel probability measures on a compact metric space is relatively compact. We shall use this fact frequently in this dissertation.

C.2 Convergence of Markov processes and operator semigroups

In this subsection, we first discuss the notion of weak convergence of a sequence of stochastic processes defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and taking values in the metric space (E, r). We then specialize to the case of Markov processes and discuss the techniques which can be employed to establish weak convergence of a sequence of Markov processes.

Since in most of our applications we will be dealing with stochastic processes having right continuous sample paths, we first study the space $D_E[0, \infty)$ of E-valued right continuous functions having left limits (RCLL) defined on $[0, \infty)$. We define the following metric on $D_E[0, \infty)$.

Definition C.2.1. *Let Λ be the collection of strictly increasing Lipshitz continuous functions λ defined on $[0, \infty)$ with $\lambda(0) = 0$ and $\lim_{t \to \infty} \lambda(t) = \infty$. For $\lambda \in \Lambda$ define*

$$\gamma(\lambda) = \sup_{0 \leq s < t} \left| \log \left(\frac{\lambda(t) - \lambda(s)}{t - s} \right) \right|. \tag{C.4}$$

Then for each $x, y \in D_E[0, \infty)$, $\lambda \in \Lambda$, and $u \in [0, \infty)$ let

$$d(x, y, \lambda, u) = \sup_{t \geq 0} r(x(t \wedge u), y(\lambda(t) \wedge u)) \wedge 1, \tag{C.5}$$

and define the metric

$$d(x, y) = \inf_{\lambda \in \Lambda} \left\{ \gamma(\lambda) \vee \int_0^\infty e^{-u} d(x, y, \lambda, u) du \right\}. \tag{C.6}$$

It can be shown that $(D_E[0, \infty), d)$ defines a metric space. The topology induced by the metric d on $D_E[0, \infty)$ is called the *Skorohod topology*. Let $\mathcal{B}(D_E[0, \infty))$ denote the σ-algebra of Borel subsets of $D_E[0, \infty)$ and $\mathcal{P}(D_E[0, \infty))$ denote the set of (Borel) probability measures defined on $(D_E[0, \infty), \mathcal{B}(D_E[0, \infty)))$.

A stochastic process X, defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and having sample paths in $D_E[0, \infty)$, is a measurable mapping $X : \Omega \to D_E[0, \infty)$. Hence, the distribution P of the process X is given by $P(\Gamma) = \mathbb{P}(X \in \Gamma)$ for all $\Gamma \in \mathcal{B}(D_E[0, \infty))$. A sequence of processes $\{X_n, n \geq 1\}$ having sample paths in $D_E[0, \infty)$ is said to converge to the process X, also having sample paths in $D_E[0, \infty)$ (written as $X_n \Rightarrow X$), if $P_n \Rightarrow P$, where for each $n \geq 1$, P_n denotes the distribution of X_n and P denotes the distribution of X. The sequence of processes $\{X_n, n \geq 1\}$ is called relatively compact if $\{P_n, n \geq 1\}$ is relatively compact. A sufficient condition for weak convergence of stochastic processes with sample paths in $D_E[0, \infty)$ is given in the following theorem.

Theorem C.2.1 ([28], p. 131, Theorem 7.8)**.** *Let E be separable and X_n, $n = 1, 2, \ldots$, and X be processes with sample paths in $D_E[0, \infty)$. If $\{X_n, n \geq 1\}$ is relatively compact and there exists a dense set $D \subset [0, \infty)$ such that*

$$(X_n(t_1), \ldots, X_n(t_k)) \Rightarrow (X(t_1), \ldots, X(t_k)), \tag{C.7}$$

for every finite set $\{t_1, \ldots, t_k\} \subset D$, then $X_n \Rightarrow X$.

So far, we have discussed some key results on the weak convergence of a sequence of stochastic processes having right continuous sample paths in a separable metric space. We now specialize on Markov processes and discuss the tools to establish weak convergence of Markov processes having values in a compact metric space. We first define transition function for a Markov processes.

Definition C.2.2. *A function $P(t, x, \Gamma)$ defined on $[0, \infty) \times E \times \mathcal{B}(E)$ is said to be a time homogeneous transition function if the following conditions are satisfied*

1. *For each* $(t, x) \in [0, \infty) \times S$*, we have* $P(t, x, \cdot) \in \mathcal{P}(E)$*, i.e.,* $P(t, x, \cdot)$ *is a Borel probability measure on* E*.*

2. *For each* $x \in E$*, we have* $P(0, x, \cdot) = \delta_x(\cdot)$*, where* δ_x *is the Dirac measure centered around* x*.*

3. *For each* $t, s \geq 0$*,* $x \in E$*, and* $\Gamma \in \mathcal{B}(E)$*, we have*

$$P(t + s, x, \Gamma) = \int P(s, y, \Gamma) P(t, x, dy). \tag{C.8}$$

A stochastic process X with state space E is said to be a *time homogeneous Markov process* with transition function $P(t, x, \Gamma)$ if for all $s, t \geq 0$ and bounded real valued Borel measurable function f on E the following holds

$$\mathbb{E}\left[f(X(t + s)|\sigma(X(u), 0 \leq u \leq t)\right] = \int f(y) P(s, X(t), dy). \tag{C.9}$$

With a time-homogeneous Markov process one can associate a group of operators satisfying the semigroup property. The precise definition is given below

Definition C.2.3. *Let* X *be a Markov process with transition function* $P(t, x, \Gamma)$*. Define an indexed family* $T = \{T(t), t \geq 0\}$ *of bounded linear operators on* $\bar{C}(E)$ *as*

$$T(t)f(x) = \int f(y) P(t, x, dy), \tag{C.10}$$

for each $f \in \bar{C}(E)$*. The family* $T = \{T(t), t \geq 0\}$ *is said to be the semigroup of operators corresponding to the Markov process* X *since it satisfies the semigroup property, i.e.,* $T(s+t) = T(s) \circ T(t)$*, where* $\circ$ *denotes composition of operators.*

Clearly, $T(0) = I$ where I denotes the identity operator on $\bar{C}(E)$. The semigroup of operators $T = \{T(t), t \geq 0\}$ is called a *contraction* semigroup if $\|T(t)f\|_\infty \leq \|f\|_\infty$ for all $t \geq 0$ and $f \in \bar{C}(E)$. Note that the semigroup $T = \{T(t), t \geq 0\}$ corresponding to the Markov process X is by definition a contraction semigroup. We also note that if $f \in \bar{C}(E)$ is such that $f \geq 0$ then by definition $T(t)f \geq 0$ for all $t \geq 0$. This property is called the positivity of the semigroup T. The semigroup $T = \{T(t), t \geq 0\}$ corresponding to the Markov process X is called Feller if

1. $\lim_{t \downarrow 0} T(t)f = f$ for all $f \in \bar{C}(E)$. (Strong continuity)

2. $T(t)1 = 1$ for all $t \geq 0$, where $1(x) = 1$ for all $x \in E$.

3. For each $t \geq 0$ and $f \in \bar{C}(E)$, we have $T(t)f \in \bar{C}(E)$.

We now state a key result which provides a sufficient condition for sequence of Markov processes to converge to a limiting Markov process in terms of their corresponding operator semigroups. We shall be using this result repeatedly in this dissertation.

Theorem C.2.2 ([28], p. 172, Theorem 2.11)**.** *Let (E, r) be a compact metric space and X be a Markov process having sample paths in $D_E[0, \infty)$ with initial distribution $\nu \in \mathcal{P}(E)$. Let $T = \{T(t), t \geq 0\}$ denote the semigroup of operators corresponding to the process X. Assume that T is Feller. For each $n \geq 1$, let X_n be a Markov process with operator semigroup $T_n = \{T_n(t), t \geq 0\}$ and having sample paths in $D_{E_n}[0, \infty)$, where $E_n \subset E$. Suppose that the following holds*

$$\lim_{n\to\infty} \sup_{x\in E_n} |T_n(t)f(x) - T(t)f(x)| = 0, \tag{C.11}$$

for each $f \in \bar{C}(E)$ and $t \geq 0$, i.e., $T_n f \to Tf$ for each $f \in \bar{C}(E)$. If $\{X_n(0), n \geq 1\}$ converges in distribution to $\nu \in \mathcal{P}(E)$, then $X_n \Rightarrow X$.

Hence, the above theorem states that a sequence of Markov processes converge to a limiting Markov process if the corresponding operator semigroups and the initial distributions converge. An effective way of establishing convergence of operator semigroups is by showing convergence of their corresponding generators which are defined below.

Definition C.2.4. *The (infinitesimal) generator of a semigroup $T = \{T(t), t \geq 0\}$ is a linear operator A on $\bar{C}(E)$ defined as*

$$Af = \lim_{t\downarrow 0} \frac{T(t)f - f}{t}, \tag{C.12}$$

for all $f \in \bar{C}(E)$ such that the above limit exists. The space on which the A is defined is called the domain $\mathcal{D}(A)$ of A. A subspace D of $\mathcal{D}(A)$ is said to be the core of A if the closure of the restriction of A to D is equal to A.

The core of the generator A of semigroup $T = \{T(t), t \geq 0\}$ can be identified using the following proposition.

Proposition C.2.1 ([28], p. 17, Proposition 3.3)**.** *Let A be the generator of a strongly continuous contraction semigroup $\{T(t), t \geq 0\}$ on $\bar{C}(E)$. Let D_0 and D be dense subspaces of $\bar{C}(E)$ with $D_0 \subset D \subset \mathcal{D}(A)$, where $\mathcal{D}(A)$ is the domain of A. If $T(t) : D_0 \to D$ for all $t \geq 0$, then D is the core of A.*

Finally, we provide the necessary and sufficient condition for convergence of a sequence of operator semigroups in terms of their corresponding generators.

Theorem C.2.3 ([28], p. 28, Theorem 6.1)**.** *For $n \in \mathbb{N}$, let T_n and T be strongly continuous contraction semigroups on $\bar{C}(E)$ with generators A_n and A, respectively. Let $D \subset \mathcal{D}(A) \subset \bar{C}(E)$ be the core of A. Then the following statements are equivalent*

(i) For each $f \in \bar{C}(E)$, $T_n(t)f \to T(t)f$ for all $t \geq 0$, uniformly on bounded intervals.

(ii) For each $f \in \bar{C}(E)$, $T_n(t)f \to T(t)f$ for all $t \geq 0$.

(iii) For each $f \in D$, $A_n f \to Af$.

References

[1] R. Schassberger, "A new approach to the $M/G/1$ processor-sharing queue," *Advances in Applied Probability*, vol. 16, no. 1, pp. 202–213, 1984.

[2] 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.

[3] M. Mitzenmacher, *The power of two choices in randomized load balancing*. PhD thesis, Harvard University, 1996.

[4] S. R. E. Turner, *Resource Pooling in Stochastic Networks*. PhD thesis, University of Cambridge, 1996.

[5] R. A. Holley and T. M. Liggett, "Ergodic theorems for weakly interacting infinite systems and the voter model," *Annals of Applied Probability*, vol. 3, no. 4, 1975.

[6] P. Clifford and A. Sudbury, "A model for spatial conflict," *Biometrika*, vol. 60, no. 3, pp. 581–588, 1973.

[7] J. T. Cox, "Coalescing random walks and voter model consensus times on the torus in $\mathbb{Z}^d$," *Annals of Probability*, vol. 17, no. 4, pp. 1333–1366, 1989.

[8] S. Galam, "Minority opinion spreading in random geometry," *European Journal of Physics*, vol. 25, no. 4, pp. 403–406, 2002.

[9] P. L. Krapivsky and S. Redner, "Dynamics of majority rule in two-state interacting spin systems," *Physical Review Letters*, vol. 90, no. 23, p. 238701, 2003.

[10] T. G. Kurtz, "Solutions of ordinary differential equations as limits of pure jump markov processes," *Journal of Applied Probability*, vol. 7, no. 1, pp. 49–58, 1970.

[11] A. S. Sznitman, "Topics in propagation of chaos," in *École d'été de probabilites de Saint-Flour XIX - 1989*, vol. 1464 of *Lecture Notes in Mathematics*, pp. 165–251, Springer, 1991.

[12] M. Kac, "Foundations of kinetic theory," in *Proceedings of the 3rd Berkeley Symposium on Mathematical Statistics and Probability*, pp. 154–161, 1956.

[13] A. L. Stolyar, "Asymptotic behavior of the stationary distribution for a closed queueing system," *Problems of Information Transmission*, vol. 25, no. 4, pp. 90–92, 1989.

[14] V. Anantharam, "A mean field limit for a lattice caricature of dynamic routing in circuit switched networks," *Annals Applied Probability*, vol. 1, no. 4, pp. 481–503, 1991.

[15] V. Anantharam and M. Benchekroun, "A technique for computing sojourn times large networks of interacting queues," *Probability in Engineering and Informational Sciences*, vol. 7, pp. 441–464, 1993.

[16] F. Baccelli, M. Lelarge, and D. R. Mcdonald, "Metastable regimes for multiplexed TCP flows," in *Proceedings of the 42nd Allerton Conference on Communication, Control, and Computing*, 2004.

[17] S. Kumar and L. Massoulié, "Integrating streaming and file-transfer internet traffic: fluid and diffusion approximations," *Queueing Systems*, vol. 55, no. 4, pp. 195–205, 2007.

[18] E. Schurman and J. Brutlag, "The user and business impact on server delays, additional bytes and http chunking in web search," in *O'Reilly Velocity Web Performance and Operations Conference*, 2009.

[19] K. Salchow, "Load balancing 101: Nuts and bolts," in *White Paper, F5 Networks, Inc.*, 2007.

[20] 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," *Performance Evaluation*, vol. 68, no. 11, pp. 1056–1071, 2011.

[21] 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.

[22] F. A. Haight, "Two queues in parallel," *Biometrika*, vol. 45, no. 3/4, pp. 401–410, 1958.

[23] J. F. C. Kingman, "Two similar queues in parallel," *Annals of Mathematical Statistics*, vol. 32, no. 4, pp. 1314–1323, 1961.

[24] A. Ephremides, P. Variya, and J. Walrand, "A simple dynamic routing problem," *IEEE Transactions on Automatic Control*, vol. 25, no. 4, pp. 690–693, 1980.

[25] R. R. Weber, "On the optimal assignment of customers to parallel servers," *Journal of Applied Probability*, vol. 15, pp. 406–413, 1978.

[26] V. Gupta, M. H. Balter, K. Sigman, and W. Whitt, "Analysis of join-the-shortest-queue routing for web server farms," *Performance Evaluation*, vol. 64, no. 9-12, pp. 1062–1081, 2007.

[27] Y. Azar, A. Broder, A. Karlin, and E. Upfal, "Balanced allocations," in *Proceedings of the 26th ACM Symposium on Theory and Computing*, pp. 593–602, 1994.

[28] S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence*. John Wiley and Sons Ltd, 1985.

[29] C. Graham, "Chaoticity on path space for a queueing network with selection of shortest queue among several," *Journal of Applied Probability*, vol. 37, no. 1, pp. 198–211, 2000.

[30] J. B. Martin and Y. M. Suhov, "Fast Jackson networks," *Annals of Applied Probability*, vol. 9, no. 3, pp. 854–870, 1999.

[31] D. Shah and B. Prabhakar, "The use of memory in dynamic load balancing," in *Proceedings of the IEEE International Symposium on Information Theory*, 2002.

[32] M. Bramson, Y. Lu, and B. Prabhakar, "Randomized load balancing with general service time distributions," in *Proceedings of the ACM SIGMETRICS*, pp. 275–286, 2010.

[33] M. Bramson, "Stability of join the shortest queue networks," *Annals of Applied Probability*, vol. 21, no. 4, pp. 1568–1625, 2011.

[34] 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.

[35] J. Xu and B. Hajek, "The supermarket game," *Stochastic Systems*, vol. 3, no. 2, pp. 405–441, 2013.

[36] E. Altman, U. Ayesta, and B. J. Prabhu, "Load balancing in processor sharing systems," *Telecommunication Systems*, vol. 47, no. 1-2, pp. 35–48, 2008.

[37] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.

[38] K. Deimling, *Ordinary Differential Equations in Banach Spaces*, vol. 96 of *Lecture Notes in Mathematics*. Springer, 1977.

[39] F. P. Kelly, *Reversibility and Stochastic Networks*. John Wiley and Sons Ltd, 1979.

[40] S. Zachary, "A note on insensitivity in stochastic networks," *Journal of Applied Probability*, vol. 44, no. 1, pp. 238–248, 2007.

[41] A. Mukhopadhyay and R. R. Mazumdar, "Randomized routing schemes for large processor sharing systems with multiple service rates," in *Proceedings of the ACM SIGMETRICS*, pp. 555–556, 2014.

[42] A. Mukhopadhyay and R. R. Mazumdar, "Rate-based randomized routing in large heterogeneous processor sharing systems," in *Proceedings of the 26th International Teletraffic Congress (ITC 26)*, pp. 1–9, 2014.

[43] 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. PP, no. 99, 2015.

[44] A. Mukhopadhyay, A. Karthik, and R. R. Mazumdar, "Randomized assignment of jobs to servers in heterogeneous clusters of shared servers for low delay." arXiv:1502.05786 [cs.DC].

[45] "Amazon EC2." http://aws.amazon.com/ec2/.

[46] "Microsoft Azure." http://www.microsoft.com/windowsazure/.

[47] V. Gupta and A. Radovanovic, "Online stochastic bin packing." arXiv:1211.2687, 2012.

[48] 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.

[49] 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.

[50] 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.

[51] C. Graham and S. Méléard, "Stochastic particle approximations for generalized Boltzmann models and convergence estimates," *The Annals of Probability*, vol. 25, no. 1, pp. 115–132, 1997.

[52] V. V. Marbukh, "Fully connected message-switching network with a large number of nodes, bypass routes and a limited number of waiting places at the nodes," *Problems of Information Transmission*, vol. 21, pp. 154–161, 1986.

[53] V. V. Marbukh, "Asymptotic investigation of a complete communications network with a large number of points and bypass routes," *Problems of Information Transmission*, vol. 16, pp. 212–216, 1982.

[54] R. J. Gibbens, P. J. Hunt, and F. P. Kelly, "Bistability in communication networks," in *DISORDER IN PHYSICAL SYSTEMS*, pp. 113–128, Oxford University Press, 1990.

[55] N. Antunes, C. Fricker, P. Robert, and D. Tibi, "Stochastic networks with multiple stable points," *The Annals of Probability*, vol. 36, no. 1, pp. 255–278, 2008.

[56] P. Nain, "Qualitative properties of the Erlang blocking model with heterogeneous user requirements," *Queueing Systems*, no. 2, pp. 189–206, 1990.

[57] M. S. Berger, *Nonlinearity and Functional Analysis: Lectures on Nonlinear Problems in Mathematical Analysis.* Academic Press, 1977.

[58] B.Kuo, *Automatic Control Systems.* Prentice-Hall, 1982.

[59] L. Markus and H. Yamabe, "Global stability criteria for differential systems," *Osaka Mathematical Journal.*, vol. 12, no. 2, pp. 305–317, 1960.

[60] J. Kaufman, "Blocking in a shared resource environment," *IEEE Transactions on Communications*, vol. 29, no. 10, pp. 1474–1481, 1981.

[61] J. W. Roberts, "A service system with heterogeneous user requirement," in *Performance of Data Communications Systems and Their Applications*, pp. 423–431, 1981.

[62] A. Mukhopadhyay, A. Karthik, R. R. Mazumdar, and F. Guillemin, "Mean field and propagation of chaos in multi-class heterogeneous loss models," *Performance Evaluation*, vol. 91, pp. 117–131, 2015.

[63] A. Mukhopadhyay, R. R. Mazumdar, and F. M. Guillemin, "The power of randomized routing in heterogeneous loss systems," in *Proceedings of the 27th International Teletraffic Congress (ITC 27)*, pp. 125–133, 2015.

[64] M. E. Yildiz, R. Pagliary, A. Ozdaglar, and A. Scaglione, "Voting models in random networks," in *Information Theory and Applications Workshop (ITA)*, pp. 1–7, 2010.

[65] P. Chen and S. Redner, "Consensus formation in multi-state majority and plurality models," *Journal of Physics A: Mathematical and General*, vol. 38, no. 33, p. 7239, 2005.

[66] P. Chen and S. Redner, "Majority rule dynamics in finite dimensions," *Physical Review E*, vol. 71, no. 3, p. 036101, 2005.

[67] F. M. Bass, "A new product growth for model consumer durables," *Management Science*, vol. 15, no. 5, pp. 215–227, 1969.

[68] N. T. J. Bailey, *The Mathematical Theory of Infectious Diseases*. Hafner, New York, 197.

[69] W. Kermack and A. McKendrick, "Contributions to the mathematical theory of epidemics: Further studies of the problem of endemicity," *Bulletin of Mathematical Biology*, vol. 53, no. 1, pp. 89–118, 1991.

[70] S. Bikchandani, D. Hirshleifer, and I. Welch, "A theory of fads, fashion, custom, and cultural change as information cascades," *Journal of Political Economy*, vol. 100, no. 5, pp. 992–1026, 1992.

[71] G. Ellison and D. Fudenberg, "Rules of thumb for social learning," *Journal of Political Economy*, vol. 101, no. 4, pp. 612–643, 1993.

[72] G. Ellison and D. Fudenberg, "Word-of-mouth communication and social learning," *The Quarterly Journal of Economics*, vol. 110, no. 1, pp. 93–125, 1995.

[73] A. Banerjee and D. Fudenberg, "Word of mouth learning," *Games and Economic Behavior*, vol. 46, no. 1, pp. 1–21, 2004.

[74] K. Chatterjee and S. H. Xu, "Technology diffusion by learning from neighbours," *Advances in Applied Probability*, vol. 36, no. 2, pp. 355–376, 2004.

[75] D. Acemoglu, M. Dahleh, I. Lobel, and A. Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.

[76] M. H. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[77] J. N. Tsitsiklis, *Problems in decentralized decision making and computation*. PhD thesis, M.I.T, 1984.

[78] P. L. Krapivsky, "Kinetics of monomer-monomer surface katalytic reactions," *Physical Review A*, vol. 45, no. 2, 1992.

[79] V. Sood and S. Redner, "Voter model on heterogeneous graphs," *Physical Review Letters*, vol. 94, no. 17, 2005.

[80] M. Mobilia, A. Petersen, and S. Redner, "On the role of zealotry in the voter model," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2007, no. 8, p. P08029, 2007.

[81] E. Yildiz, A. Ozdaglar, D. Acemoglu, A. Saberi, and A. Scaglione, "Binary opinion dynamics with stubborn agents," *ACM Transactions on Economics and Computation*, vol. 1, no. 4, pp. 1–30, 2013.

[82] E. Mossel, J. Neeman, and O. Tamuz, "Majority dynamics and aggregation of information in social networks," *Autonomous Agents and Multi-Agent Systems*, vol. 28, no. 3, pp. 408–429, 2014.

[83] P. Flocchini, E. Lodi, F. Luccio, L. Pagli, and N. Santoro, "Dynamic monopolies in tori," *Discrete Applied Mathematics*, vol. 137, no. 2, pp. 197–212, 2004.

[84] G. Moran, "The r-majority vote action on 0-1 sequences," *Discrete Mathematics*, vol. 132, no. 1, pp. 145 – 174, 1994.

[85] Z. Agur, A. Fraenkel, and S. Klein, "The number of fixed points of the majority rule," *Discrete Mathematics*, vol. 70, no. 3, pp. 295 – 302, 1988.