BACHELOR

Mean Field Approximations for Heterogeneous Supermarket Models on Graphs

Devlin, Finley

Award date:
2022

Department of Mathematics and Computer Science

Mean Field Approximations for Heterogeneous Supermarket Models on Graphs

Bachelor Thesis

Finley Devlin

Supervisors:
Prof.dr.ir S.C. Borst
E. Cardinaels

Eindhoven, July 2022

Abstract

In this paper we consider a variation of the so called supermarket model with heterogeneous processing rates and underlying network structure that cause restrictions in load balancing. Servers exist on the vertices of a graph where an edge denotes communication between servers. Jobs arrive according to a Poisson process and are assigned by a dispatcher to a server with exponentially distributed processing times of heterogeneous rate. The processing rates may be different for each server thus The dispatcher randomly samples two servers connected in the graph and assigns the job to the shorter of the two queues with ties broken uniformly. This is the well known power-of-two policy applied on a graph topology.

We take a mean field approach based on the generalised mean field approximation for heterogeneous objects by N. Gast and S. Allmeier (2022). We review existing mean field methods for homogeneous systems and compare the interpretations between the two. To this end, there is difficulty in establishing a scaling limit interpretation of the heterogeneous mean field for our model given the lack of a 'natural' scaling scheme for the heterogeneous model on graphs.

We suggest the following scaling regime for a system with arbitrary service rates and graph structure: replace each server by a cluster of servers all with processing rate equal to the replaced server and all clusters the same size. Then say that two servers are connected if their clusters are connected. We find that the mean field is equal for a server in the original system and any server in the corresponding cluster for any degree of scaling. Thus suggesting cluster scaling preserves the heterogeneity of the system. Following this we show the limit of the expected occupancy state for a given cluster k is equal to the mean field of corresponding server k from the original system. As such the heterogeneous mean field may be interpreted as an approximation for a finite system or the limit of the expected occupancy state for a sequence of cluster scaled systems.

Contents

Chapter 1

Introduction

This work explores the mean field approximation and its application to heterogeneous load balancing systems. We consider a variation of the so called supermarket model with heterogeneous processing rates and underlying network structure that cause restrictions in load balancing. In this model jobs arrive according to a Poisson process and are assigned by a dispatcher to a server with exponentially distributed processing times of heterogeneous rate. The assignment is determined by a load balancing policy which is restricted by the network structure represented by a suitable graph. The specific policy under investigation is the so called Join-the-Shortest-Queue-2 (JSQ(2)) policy, also known as the power-of-two policy, where an arriving job is assigned to the shortest of two randomly sampled queues. In our model, the network restrictions mean the two sampled queues must be connected in the underlying graph.

In recent times, load balancing problems have become more prevalent with the advent of large scale data centres and cloud computing networks. Our models are of particular relevance to cloud computing where there are often many different types of servers which do not necessarily communicate between one another. Heterogeneous processing rates model different types of servers with different processor speeds and the underlying network structure models communication channels such that servers can share jobs between them.

The mean field approximation is useful in large scale load balancing problems where analytic distributions of queue lengths can easily become intractable. The main idea behind the approximation is to model the evolution of the network completely by the expected change in the network. This gives a system of ordinary differential equation that we solve to obtain an approximation of the queue length distribution known as the mean field. The mean field approximation has long been considered for analysis of homogeneous systems appearing in classical work [1, 2]. However, only recently in [3] has it been generalised for heterogeneous systems in which exchangeability between servers is not present.

Given the increasingly large size of data centres and cloud computing networks, the scalability of load balancing policies is important for analysis of their performance. In this work we pay particular interest to the growth of our chosen queueing networks and its effect on the error between the queue length distribution and mean field approximation. A challenge for heterogeneous systems is that it is not always clear how they scale. Despite this, in [3] it was shown that even in heterogeneous systems the error between the queue length distribution and mean field approximation is $O(1/n)$ with n the network size and thus vanishes as the networks grow large.

The overarching goal of the work is to study the application of the 'new' mean field approximation from [3] to heterogeneous queueing networks on graphs. We discuss how to interpret this 'new' approximation and attempt to expand this through comparison with classical mean field work on homogeneous systems. Furthermore, we discuss issues with scaling heterogeneous systems and suggest a type of scaling that preserves their structure and heterogeneity. We also investigate the aforementioned $O(1/n)$ error term from [3, Thereom 4.1] and its implications for scaling heterogeneous systems.

In Chapter 2 we give an overview of some relevant work on the the mean field approximation and load balancing with underlying network structures. We also describe the model and give insight into different interpretations of the mean field. We begin our analysis in Chapter 3 with models of heterogeneous processing rates but no underlying structure. Conversely, Chapter 4 looks at homogeneous systems with underlying graph structures. Through this approach we study the separate impact of the graph structure and heterogeneous rates. We then apply the concepts from Chapters 3 and 4 in Chapter 5 where we consider systems with both heterogeneous processing rates and underlying graph structures. We discuss how these systems scale and attempt to generalise results from the previous two chapters. Finally, our main conclusions are summarised in Chapter 6.

Chapter 2

Background

2.1 Related Literature

The JSQ(2) policy which we use in this work is a special case with $d = 2$ of the Join-the-Shortest-Queue-d or power-of-d policy which was first studied in a queueing context by Mitzenmacher in [1] and Vvedenskaya *et al.* in [2]. The d refers to sampling d many servers upon a job arrival and assigning it to the shortest of the d queues. This policy was applied to the so called supermarket model which is a queueing network of n servers, each with exponential processing times with homogeneous rate μ and Poisson arrivals at rate $n\lambda$. Arrivals are then assigned to one of the n queues at arrival by the dispatcher according to the policy. An extensive overview of the literature on power-of-d policies and other scalable load balancing issues can be found in [4].

Power-of-d policies can become quickly intractable which motivates the use of approximations and asymptotic methods to analyse the performance of these policies. One such approximation is known as the mean field approximation which is a well known method from Interacting Particle Systems and Statistical Mechanics [5]. Mean field methods see wide spread use in load balancing problems [6]. Mitzenmacher's analysis of the homogeneous supermarket model in [1] is often referred to as a mean field method which is explained in Section 2.3.1. Mean field methods have also been used in the analysis of other load balancing policies for specific heterogeneous systems [7, 8, 9].

We pay particular attention to the recent work [3] which generalises the mean field approximation for general heterogeneous systems. A more detailed overview of this approximation which we dub the 'new mean field' will be given in Section 2.3.2. A particularly important aspect of [3] is that it establishes an accuracy of the mean field which we will also later discuss. Furthermore, in [3] they consider the 'refined mean field' which is an adjusted version of the mean field approximation that is more accurate. They expand the existing work on the refined mean field [6, 10] to heterogeneous systems, although we focus on the unrefined mean field.

In our work we study load balancing for systems where choices are restricted by a graph structure. An overview of load balancing on graphs is outlined in [4]. Gast studied the mean field method for the power-of-two policy on graphs in [11] which was found to perform poorly as an approximation. This motivated a new approximation called the 'pair approximation' developed by Gast which proved to be more accurate than the mean field for graphs although more specialised. Other work regarding systems on graphs focuses on the asymptotic behaviour of system, examples of which can be found in [12, 13, 14]. The model used in [12] is nearly identical to our own and we compare our results in Section 3.3.2. They take a rigorous fluid limit approach to constructing the mean field for homogeneous systems on graphs which will be discussed in Section 2.3.1.

Finally, we note that the literature for systems with both heterogeneous processing rates and underlying graph structures is sparse. In [15] Ramanan presents the open problem to study the performance of the JSQ(2) policy with general service distributions and underling graph structures. The article [15] is included in the special issue '100 views on queues' Volume 100 of the journal Queueing Systems.

2.2 Model Overview

Throughout this work we consider variations on the model known as the supermarket model on graphs also seen in [12]. Consider a network with n servers situated on the vertices of a graph $G_n = (V_n, E_n)$ where $V_n = \{1, ..., n\}$ unless stated otherwise. We have Poisson arrivals with rate λn to a dispatcher which upon an arrival uniformly chooses a server k from any in the network. Then a neighbour of k is chosen uniformly from its interaction neighbourhood (Definition 1). The arrival is assigned to the shorter of the two queues with ties broken uniformly. Furthermore, each server $k \in V_n$ has exponential processing times with rate μ_k. We consider this policy applying JSQ(2) to some graph topology where the second choice is restricted by the local neighbourhood. A graphical representation of the model can be seen in Figure 2.1.

Definition 1 (Interaction neighbourhood)**.** The (interaction) neighbourhood of a server k, node of graph $G_n = (V_n, E_n)$, is given by

$$\mathcal{N}_k^{(n)} := \{l \in V_n \ : \ (k, l) \in E_n\} \quad \text{for all } \ k \in V_n.$$

An alternative interpretation of this model is to observe that there exists an edge between any two sampled servers. Therefore, we can see this as arriving to a random edge and then choosing the shorter queue of the two nodes. Both interpretations have physical analogues. The former can be imagined in a cloud computing network where there is no centralised dispatcher. In this setting jobs arrive to a random or 'nearest' server (assuming jobs are distributed randomly) and then there is some local dispatcher at each server that samples a second server and assigns the arrival. The latter interpretation can be seen as a network with a single distributor, it randomly samples an edge, requests the queue lengths and then assigns the arrival to the shortest queue. This idea of arrivals to an edge is somewhat unique to JSQ(2), since for JSQ(d), $d \geq 3$ it would require sampling multiple edges connected to one point, or seen as sampling a subgraph.

It is important to note that with this model we cannot sample the same server twice because it does not belong to its own interaction neighbourhood (Definition 1). This means that even when using the fully connected graph, this is not the same model that appears in [1, 2, 3] since these models sample servers with replacement. We will notice in Section 4.4 that in the

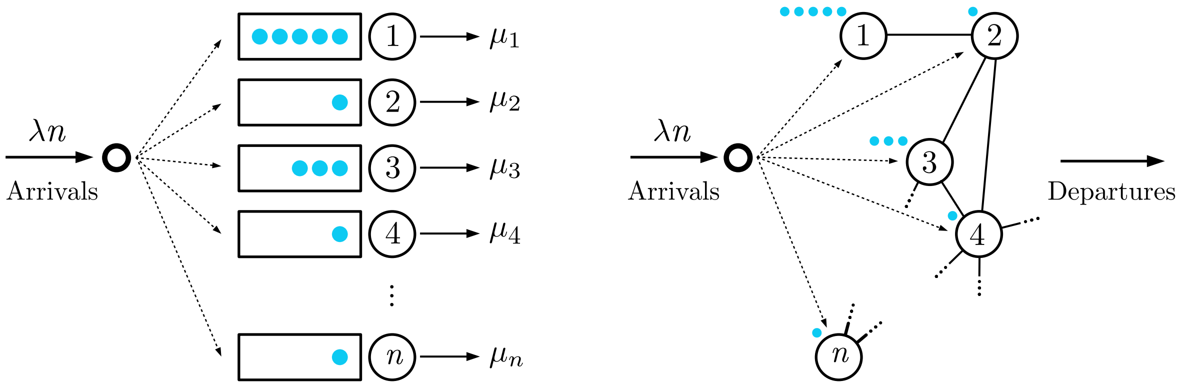

Figure 2.1: Overview of the model with Poisson arrivals at rate λn to a dispatcher that assigns the job to one of n servers each with heterogeneous processing rate μ_k according to JSQ(2) on graphs. The heterogeneity of the system is demonstrated on the left and on the right there is an example underlying graph structure.

homogeneous case this makes no difference to the mean field. However, in the heterogeneous case this does have an impact which we will explore in Sections 3 and 5. When we use the policy with replacement we define the interaction neighbourhood by $\mathcal{N}'^{(n)}_k := \mathcal{N}^{(n)}_k \cup \{k\}$. This could also be seen as a graph with a loop for each node. We interpret this policy as an arrival to a neighbourhood which is a subgraph and then sampling the second server from the vertices of the subgraph $\mathcal{N}'^{(n)}_k$.

2.3 Mean field

In this section we provide preliminaries and intuition on various forms of the mean field approximation for load balancing. Here we also highlight the interpretations for different constructions and uses of the mean field.

2.3.1 Classical mean field methods

There are three techniques that we refer to as the mean field. They are used to analyse the queue length distribution of the homogeneous supermarket model on the fully connected graph. All three yield the same result and so we refer to them collectively as the classical mean field, when in fact only one of the three is a mean field approximation in construction. Later we discuss how these interpretations apply to the new mean field on the heterogeneous supermarket model. In this section we consider the homogeneous supermarket model on the complete graph.

Mean field approximation

The idea behind this approximation is to first consider the evolution of a single server as a birth death process. Then all servers have the same processing rate and the graph is fully connected, so we say all servers behave the same in expectation. Following this, we say the rate of change is equal to the expected change in the system and we solve this system of

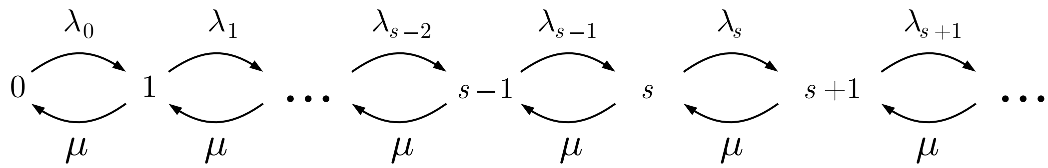

Figure 2.2: Transition rate diagram for the birth death process of the evolution of one server with state space of queue length.

ODEs to obtain our mean field approximation. An example of this method can be found in Section 2.1 of [11]. Consider the birth death process of a single server called k with state space of the queue length $\{0, 1, 2, ...\}$. The transition rate diagram can be seen in Figure 2.2. For ease of analysis we define $q_s(t)$ to be the proportion of servers that have queue length of s or greater at time t, although we drop notational dependence on t when otherwise clear. Then we can say that

$$\lambda_s = n\lambda p_s := n\lambda \cdot \mathbb{P}\left(\text{Arrival to server } k \mid \text{Server } k \text{ has queue length } s\right)$$

Upon arrival we sample two servers with replacement, we can sample server k first or second and so it is sampled with probability $2/n$. Next we consider the queue length of the other sampled server. If the queue length is greater than s the job is assigned to server k, if the queue length is s then it is assigned to server k with probability $1/2$ and otherwise it is assigned to the other server. Hence we have

$$\lambda_s = n\lambda p_s = n\lambda \cdot \frac{2}{n}\left(q_{s+1} + \frac{1}{2} \cdot (q_s - q_{s+1})\right) = \lambda(q_s + q_{s+1}).$$

Now we split the state space in two and look at the expected flow between two halves where the server k has queue length greater or equal to s. This is equal to the expected flow between states $s-1$ and s. The probability that server k is in state s is equal to $q_s - q_{s+1}$ and so for $s \geq 1$ we find the expected flow

$$\mu(q_s - q_{s+1}) + (q_{s-1} - q_s)\lambda_{s-1} = \lambda(q_s + q_{s+1})(q_{s-1} - q_s) - \mu(q_{s+1} - q_s) = \lambda(q_{s-1}^2 - q_s^2) - \mu(q_{s+1} - q_s).$$

Let $Q_s = nq_s$ give the number of servers that have queue length greater or equal to s. We approximate the rate of change of Q_s by the sum of the expected flow rates of all the servers. This sum represents the expected rate at which servers enter or leave the group of server with queue length greater or equal to s. By exchangeability of all servers we assume that they all have the same expected flow rate and so we obtain

$$\frac{dQ_s}{dt} = n\left[\lambda(q_{s-1}^2 - q_s^2) - \mu(q_{s+1} - q_s)\right], \qquad s \geq 1$$

thus division by n yields

$$\frac{dq_s}{dt} = \lambda(q_{s-1}^2 - q_s^2) - \mu(q_{s+1} - q_s), \qquad s \geq 1 \tag{2.1}$$

with $q_0 = 1$. The solution to system of ODEs (2.1) given by $q = (q_s)_{s\in\mathbb{N}}$ is what we call the mean field. Throughout its construction we interpreted q_s as a probability now we interpret q as an approximation of the queue length distribution.

A fact that will not prove true for the remaining methods is that we can find the mean field approximation for a specific n sized system. This will be important in interpreting how we use the mean field. It should also be clear that this is an approximation and provides no direct bounds on the error or metric on the accuracy.

Population process approach

This approach was taken by Mitzenmacher in [1] and uses results on a class of models known as density dependent population processes studied in [16]. This approach uses a state space of population densities

$$q^{(n)} = (q_0^{(n)}, q_1^{(n)}, ...) \in \{(Q_0/n, Q_1/n, Q_2/n, ...) \mid Q_s \in \{0, 1, ..., n\}\}$$

where Q_s is the number of servers with s or more jobs in their queue. We notice that $q_s^{(n)} \in [0, 1]$ for all n and so we drop the dependence on n inline with notation from [1]. These are known as population densities because $q_s = Q_s/n$ is the proportion or density of servers with s or more jobs in their queue. We can completely describe the system with this state space since the homogeneous servers are exchangeable. This is just a population process, we must show it is density dependent.

Let $e_s \in l^1$ be the vector with a 1 in the s index and zero elsewhere. We find the following transition rates

$$q \to q + e_s/n - e_{s-1}/n \quad \text{at rate} \quad n\lambda(q_{s-1}^2 - q_s^2) \tag{2.2}$$
$$q \to q - e_s/n + e_{s-1}/n \quad \text{at rate} \quad \mu(q_s - q_{s+1})n. \tag{2.3}$$

Let us justify these transition rates. The probability that an arriving job enters a server of length greater than or equal to s is the probability to sample two servers of queue length s or greater, which is given by q_s^2. Therefore we have arrivals to queues of length $s-1$ at rate $n\lambda(q_{s-1}^2 - q_s^2)$, yielding (2.2). Likewise jobs leave servers of queue length s at rate $\mu n(q_s - q_{s+1})$ since we have $n(q_s - q_{s+1})$ many servers of length s with processing rate μ, yielding rate (2.3).

Following reasoning from [1], the transition rates of this process only depend on the population densities q_s and so we specifically have a density dependent population process. Following Kurtz' Theorem, Mitzenmacher explains that the changes in population density become continuous as $n \to \infty$ and that the limiting system is deterministic. From this continuity we know that the rate of change of the population density q_s is equal to its expected change in q_s yielding deterministic ODE (2.4) with $q_0 = 1$. Intuitively from (2.2), (2.3) we see changes in the state becomes infinitesimal and so there is a meaning to taking the derivative of the state with respect to time.

$$\frac{dq_s}{dt} = n\lambda(q_{s-1}^2 - q_s^2) \cdot 1/n + n\mu(q_{s+1} - q_s) \cdot (-1/n), \qquad s \geq 1 \tag{2.4}$$
$$= \lambda(q_{s-1}^2 - q_s^2) - \mu(q_{s+1} - q_s), \qquad s \geq 1. \tag{2.5}$$

We clearly notice that this is the same as the ODE (2.1) and so the mean field approximation is equivalent to the density dependent population process as $n \to \infty$. As we alluded to earlier, this approach only has meaning as the deterministic limit of a sequence of population processes. Therefore we cannot take this approach to a system with fixed n whereas we are able to with the mean field approximation. This will prove to be important understanding for our remaining work. It should also be clear that this is a more rigorous approach to the mean field method. As $n \to \infty$ the effect of a change in one server on a given other server tends to zero because the change of a finite number of servers is 'drowned out' by the vastness of infinite servers. Thus the servers are said to behave independently as $n \to \infty$.

Scaling limit approach

The final and most rigorous approach is to consider the fluid or scaling limit of the model which we take from [4]. This means precisely finding the behaviour as the number of servers tend to infinity. We use a state space

$$Q(t) = (Q_0(t), Q_1(t), ...) \in \{(Q_0(t), Q_1(t), ...) \mid Q_s(t) \in \mathbb{N}\}$$

where $Q_s(t)$ is the number of servers with s or more jobs in their queue. In order to find the fluid limit, we must consider sequences of systems indexed by their size n. We define $q_s^{(n)} := Q_s^{(n)}/n$ which represents the fraction of servers in the $n-$th system with s or more jobs in their queue. Then we define the fluid limit $q = (q_0, q_1, ...)$ by a limit in distribution of each element

$$q_s^{(n)} = Q_s^{(n)}/n \to q_s \quad \text{as } n \to \infty. \tag{2.6}$$

It can then be said that the sequence of processes $\left(q^{(n)}\right)_{n\in\mathbb{N}}$ has weak limit q which satisfies the system of ODEs (2.1). The proof of this describes the process $q^{(n)}$ as a Poisson process, then uses martingales to argue for existence of the limit. Then the fluid limit can be described in terms of an integral equation. An exmaple of this approach can be seen in [12] and is less intuitive than the previous two approaches. In general, finding the fluid limit is not a trivial problem and is outside the scope of this Bachelor project.

In many cases the system of ODEs representing the fluid limit (in our case (2.1)) cannot be solved exactly but we may be able to find a fixed point when $dq_s/dt = 0$. This is recognisable as the stationary distribution, $t \to \infty$, of the fluid limit. This kind of analysis is visualised in Figure 2.3. In [1], Mitzenmacher found the stationary point $q_s = \lambda^{2^s-1}$ and provided results on the rate of convergence and stability.

In general, finding the stationary distribution of a finite system using JSQ(2) is intractable. In Queueing Systems the stationary distribution is a desirable entity for performance analysis of the system and allows for the construction of a multitude of performance metrics. Therefore, a natural question is whether we can use the stationary distribution of the fluid limit to approximate the stationary distribution of a finite system. A first step to answering this is to consider whether the stationary distribution of the fluid limit (fixed point) is equal to the fluid limit of stationary distributions. This equates to whether we can interchange the order of limits $t \to \infty, n \to \infty$. In Figure 2.3 we can see this as either taking the top or bottom path respectively from the stochastic process to the fixed point. It is also in general very difficult to prove whether the limits are interchangeable.

Similarly to the population process approach, it is clear that this approach requires taking the limit of a sequence of systems. Thus the fluid limit depends on how the systems scale.

$$\begin{tikzcd}
\text{Stochastic process} \arrow[r, "n \to \infty"] \arrow[d, "t \to \infty"'] & \text{Fluid limit} \arrow[d, "t \to \infty"] \\
\begin{matrix}\text{Stationary}\\ \text{distribution}\end{matrix} \arrow[r, "n \to \infty"] & \text{Fixed point}
\end{tikzcd}$$

Figure 2.3: Diagram representing fluid limit analysis.

There is no way to find a fluid limit for just one finite system. By showing the fluid limit is in fact the limit of the process we can have a level of confidence that it is accurate for large systems. How accurate is tackled by the new mean field in [3].

2.3.2 New Mean Field

We define the new mean field as the one described in [3] which is applicable for heterogeneous systems. This follows a similar particle like approach to the classical mean field in Section 2.3.1 but instead of considering just one server, we consider each server to allow for heterogeneity. In order to look at each server and their queue length independently we use a state space of binary objects. Consider a system with n servers each with buffer of size b then for any server $k \in V, s \in \mathcal{S} = \{1, 2, ..., b\}$ we define

$$X^{(n)}_{(k,s)}(t) := \begin{cases} 1 & \text{If server } k \text{ has queue length } s, \\ 0 & \text{Otherwise.} \end{cases} \tag{2.7}$$

Note that in our model from Section 2.2 we do not have a buffer. A buffer is necessary to have a finite state space which is required in [3] however we can in principle choose a sufficiently large finite buffer such that the dropping probability is negligible. Henceforth it is assumed that the buffer for each server is the same and sufficiently large. The state of the whole system can be described by a matrix of binary objects (2.7). We note that each row will have a singular 1 since the row states are mutually exclusive. The state at time t is given by

$$X^{(n)}(t) := \left(X^{(n)}_{(k,s)(t)}\right)_{k \in V, s \in S} \tag{2.8}$$

and thus the stochastic process is given by

$$X^{(n)} = \left\{X^{(n)}(t)\right\}_{t \geq 0}. \tag{2.9}$$

However, when referring to elements of the matrix, we often we drop the dependence on t and simply write $X^{(n)}_{(k,s)} = X^{(n)}_{(k,s)}(t)$ for ease of notation. The stochastic process $X^{(n)}$ is a continuous time Markov process because we have a Poisson arrival process and exponential service times which are both memoryless. Furthermore, it is time homogeneous since the transition rates depend on the state and not the time. We refer to the entire state space by $\mathcal{X}^{(n)} \subset \{0,1\}^{n \times |S|}$ and let $e^{(n)}_{k,s}$ be the matrix that is 1 in element (k, s) and zero elsewhere. The index n gives the number of servers in the system. We also use this index to construct sequences of systems $X = (X^{(d(n))})_{n\in\mathbb{N}}$ with $d(n)$ giving the number of servers in the n-th system.

The assumption of independent evolution of objects results in approximating each server's evolution with its expected behaviour across interactions with any other server. This is captured in the drift, which is the expected change in the system across an infinitesimal change in time. Since $X^{(n)}$ is a continuous time Markov process we can find the drift with relative ease. For the Markov process, the expected change in state in infinitesimal time over the length of time is the sum of the product of transition rates and the change in state. The following definition is taken from [3].

Definition 2. For a given state $x \in \mathcal{X}^{(n)} \subset \{0,1\}^{n\times|S|}$ the drift of a time-homogeneous Markov process $X^{(n)} = \left\{X^{(n)}(t)\right\}_{t\geq 0}$ defined in (2.9) is a function $f : \mathcal{X}^{(n)} \to \mathbb{R}^{n\times|S|}$ such that

$$f^{(n)}(x) = \lim_{\delta\to 0} \frac{1}{\delta}\mathbb{E}\left[X^{(n)}(t+\delta) - X^{(n)}(t) \mid X^{(n)}(t) = x\right]. \tag{2.10}$$

Next we construct the mean field. The idea is to replace the binary objects with continuous ones in $[0,1]$, then we approximate the rate of change in time of these new states with the drift of the continuous objects. To do this we extend the definition of the drift (Definition 2) to allow for continuous input $f^{(n)} : \mathrm{conv}(\mathcal{X}^{(n)}) \to \mathbb{R}^{n\times|S|}$. Note that the drift is still constructed from the binary objects but now allows inputs from the convex hull of the process' state space.

Definition 3. Consider a time homogeneous continuous time Markov process $X^{(n)} = \left\{X^{(n)}(t)\right\}_{t\geq 0}$ with state space $\mathcal{X}^{(n)} \subset \{0,1\}^{n\times|S|}$ and extended drift $f^{(n)}$. Then the mean field is a function $\phi^{(n)} : \mathrm{conv}(\mathcal{X}^{(n)}) \times \mathbb{R}^+ \to \mathrm{conv}(\mathcal{X}^{(n)})$ given by the solution of the ODE

$$\frac{d\phi^{(n)}(x,t)}{dt} = f^{(n)}(\phi^{(n)}(x,t)) \tag{2.11}$$

where $x \in \mathrm{conv}(\mathcal{X}^{(n)})$ is some initial condition, $x = X^{(n)}(0)$ which unless stated otherwise is the empty system.

We also provide a technical lemma on the uniqueness of the solution of (2.11), it follows from standard ODE results but a brief a proof can be found in Appendix A.1.

Lemma 1. *If $f^{(n)}_{(k,s)}$ is a polynomial for all $k \in V_n, s \in \mathcal{S}$, then the solution of (2.11) is unique.*

For a given server k the sum of $\phi^{(n)}_{(k,s)}$ across all queue lengths s is equal to 1 and they each have value in $[0,1]$ thus we may interpret this as a probability. The mean field is an approximation of the queue length distribution for each server, although this will be formalised in Theorem 2. We do not have a general solution of (2.11) and so we compute this numerically. We discuss a small portion of numerical results in Section 3.4, 4.6 and 5.5 but note that numerical considerations are well covered for the new mean field and refined mean field in [17].

Through this construction of the mean field [3] obtained bounds on the error between $\phi^{(n)}$ and the true queue length distribution. Consider a sequence of systems $X = \left(X^{(d(n))}\right)_{n\in\mathbb{N}}$ and so the n-th system is of size $d(n)$. Let $\phi^{(d(n))}$ denote the mean field for the n-th system in the sequence, then the following theorem is obtained from [3, Theorem 4.1] and adapted for our model

Theorem 2 ([3, Theorem 4.1])**.** *Consider the model as described in Section 2.2 and the above sequence. Let* $\phi^{(d(n))}(x,t)$ *be the solution of* (2.11) *with drift* $f^{(d(n))}$ *and initial condition* $x \in conv(\mathcal{X}^{(d(n))})$ *for the n-th system of the sequence. If the processing rates* μ_k *are* $O(1)$ *bounded and the arrival rates to each server pair are* $O(1/d(n))$ *bounded then we have for* $(k,s) \in V_{d(n)} \times S$ *and* $t < \infty$

$$\mathbb{P}\left(S_k^{(d(n))}(t) = s\right) = \mathbb{E}\left[X_{(k,s)}^{(d(n))}(t)\right] = \phi_{(k,s)}^{(d(n))}(x,t) + O(1/d(n)),$$

where $S_k^{(d(n))}(t)$ *is the random variable which gives the queue length of server* k *at time* t *of the* n*-th system in the sequence.*

A notable difference in notation between our work and [3] is that we explicitly refer to a sequence of models and use $d(n)$ to denote the number of servers in the n-th system. The conditions of the theorem regarding bounded processing and arrival rates refer to all of the rates for all of the systems in the sequence. The arrival rate depends on the system size and so the condition on the arrival rate requires that the arrivals to each server is bounded by the inverse of the system size.

The way to interpret Theorem 2 is that the queue length distribution of server k in the n-th system of the sequence can be approximated by $\phi_{(k,\cdot)}^{(d(n))}$ with an error term $O(1/d(n))$. The error term being of order $1/d(n)$ seems to imply that the error term vanishes as the systems grow large. However, the constant associated with the error term depends on the rates and graph structure of the system as well as the specific server, queue length and time. If we want to study the size of the error term as the sequence grows with n, then the constant must hold for all error terms in all systems in the sequence. Thus the error term $O(1/d(n))$ is interpreted by considering a constant which is the supremum of all the individual constants for all systems in the sequence. Then with this supremum constant we can compare the error terms between different systems in the sequence. Further notice that the error term vanishes more rapidly when the number of servers in the system grows more rapidly. Thus the growth of the sequence is important in determining the rate at which the error term vanishes. Furthermore, the error term does not necessarily uniformly behave in an $O(1/n)$ sense since the error term depends on the rates and structure of the system. However, we know that there exists a constant k^* (the supremum of all the constants) such that every error term is bounded by k^*/n.

Now we can see that the error term vanishes as the systems grow large and so it is temping to interpret Theorem 2 as a kind convergence of the sequence. This is not correct in general. Notice that the state space grows with $d(n)$ and so there is no way to directly compare different sized systems to obtain a convergence in distribution type result. As will be discussed in Section 3.3.2, it may be possible to construct suitable quantities such that scaling limit like results (Section 2.3.1) can be obtained.

The utility of defining sequences of systems is to define the growth of the number of servers and thus the behaviour of the error term despite the lack of (meaning to) convergence along the sequence. Note that the systems in the sequences do not necessarily behave similarly - they could for example have vastly different graph structures or processing rates. Thus the error term can vanish, but the systems do not necessarily behave similarly. Although in most cases we define how the systems grow along the sequence.

Finally, we note that the mean field itself is independent of the sequence chosen. For any fixed sized finite system we can find the new mean field. However, we are unable to make statements about, or bound, the error term without considering a sequence of other systems which is the purpose of Theorem 2.

2.3.3 Interpretation

Whilst in Section 2.3.1 we saw that the classical mean field has three interpretations, the new mean field in Section 2.3.2 only has one. Their shared interpretation is the Interacting Particle Systems type approach from Section 2.3.1. In this interpretation the mean field is an approximation of the queue length distribution where there exists some error term between the two. With this approach we can find the mean field of a singular finite system. We henceforth refer to this as the 'particle-like' interpretation of the mean field.

On the other hand, it is not clear whether there exists an interpretation of the new mean field as a population process or scaling limit as per the classical mean field in Section 2.3.1. In these approaches we interpret the mean field as the exact deterministic behaviour of the system as the number of servers $n \to \infty$. Thus the scaling limit or population process interpretation does not exist without a sequence of systems. This is starkly different to the particle like approach which we can find for a singular finite system, as used to construct the new mean field. We henceforth refer to the asymptotically exact methods as 'scaling-like' interpretations of the mean field.

To understand how these two interpretations coexist, consider a sequence of systems whose size grows to infinity. We can find the particle like mean field of any element of the sequence since this is a finite system. We may also be able to find a scaling like mean field for the limit of the sequence. Then depending on the sequence, these may or may not be the same. For the homogeneous system on a fully connected graph, they are the same.

As mentioned in Section 2.3.2, the difficulty in constructing limiting statements of the new mean field is that the state space grows with the system size. Therefore if we want to make limiting statements we must be able to describe the state space by quantities that do not change size with n. This could entail aggregating like (exchangeable) servers and looking at the proportions of like servers. However, with heterogeneous service rates and graph structures, it is also not immediately clear what the conditions are for servers to behave similarly or be exchangeable. Furthermore, any attempt to construct a scaling like mean field would require a sequence of systems with well defined growth which is not necessarily obvious for heterogeneous systems.

Now consider that in the homogeneous fully connected system, both the particle and scaling interpretations yield the same mean field. Thus the mean field of a given finite system could be seen as the the particle like mean field or as the scaling limit of a sequence of systems starting from the given finite system. Furthermore, there is a natural way to scale the homogeneous model on the complete graph - just maintain the same homogeneous rate and use the complete graph for all system sizes. However, for a finite heterogeneous model, there exist many more ways to scale the model. A question that follows is can we find a way to scale a given finite heterogeneous system such that its particle like mean field is equal to a scaling like mean field of a sequence of scaled systems?

Chapter 3

Heterogeneous service rates

We first acquaint ourselves with the new mean field on the fully connected graph with heterogeneous service rates. Using the fully connected graph allows us to focus on the effect of the service rates.

3.1 Model

We use the model outlined in Section 2.2. For a system of n servers we have Poisson arrivals at rate λn and exponential processing times with heterogeneous processing rate μ_k for server k. We then apply JSQ(2) on the fully connected graph with replacement and so we are able to choose the same server twice. Our justification is that this allows us to draw parallels with classical work [1, 2]. The stability condition is analysed in [18], for the purposes of our work we assume the system is stable.

We use notation from Section 2.3.2 and so the state space is given by the $n \times b$ matrix of binary objects $X^{(n)} = X^{(n)}(t)$ (we drop the dependence on t for ease of notation). We first describe the transition rates of the model. For any two servers $1 \leq k, k_1 \leq n$ and queue lengths $s, s_1 \in S$ of servers k, k_1 respectively, we have partial transition rates

$$X^{(n)} \to X^{(n)} - e^{(n)}_{(k,s)} + e^{(n)}_{(k,s-1)} \quad \text{at rate } \mu_k X^{(n)}_{(k,s)} \mathbb{1}_{\{s>0\}} \tag{3.1}$$

$$X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)} \quad \text{at rate } (2\lambda \mathbb{1}_{\{s_1>s\}} + \lambda \mathbb{1}_{\{s_1=s\}}) X^{(n)}_{(k,s)} \frac{X^{(n)}_{(k_1,s_1)}}{n}. \tag{3.2}$$

Notice that the binary objects $X^{(n)}_{(k,s)}$ act as indicators that server k is of length s. The first rates (3.1) are straightforward. If server k is of queue length s then it finishes a job and has queue length $s-1$ with its processing rate μ_k assuming that $s > 0$, otherwise if $s = 0$ then there is no job to finish. For the second rates (3.2), which represent arrivals, we first uniformly choose server k of some queue length s with probability $X^{(n)}_{(k,s)}/n$ with replacement. We next uniformly choose another server k_1 of queue length s_1 also with probability $X^{(n)}_{(k_1,s_1)}/n$. By JSQ(2), if $s_1 > s$ then there is an arrival rate of λn to server k, if $s = s_1$ then they are split uniformly so we have rate $\lambda n/2$ by splitting the Poisson Process [19] and if $s > s_1$ then there is a rate of zero. Summing these three cases we multiply them by the probability to choose the configuration of servers. We finally note that we can instead choose server k_1 first and then server k which gives us the factor of 2 found in (3.2). Then the full transition rates for the arrivals are found by fixing (k, s) and summing over all possible (k_1, s_1). Hence the rate $X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)}$ is given by

$$\begin{aligned}\sum_{1\leq k_1\leq n,\ s_1\in S} (2\lambda \mathbb{1}_{\{s_1>s\}} + \lambda \mathbb{1}_{\{s_1=s\}}) X_{(k,s)} \frac{X_{(k_1,s_1)}}{n} &= \lambda X_{(k,s)} \sum_{1\leq k_1\leq n,\ s_1\in S} (2\mathbb{1}_{\{s_1>s\}} + \mathbb{1}_{\{s_1=s\}}) \frac{X_{(k_1,s_1)}}{n} \\ &= \lambda X_{(k,s)} \sum_{1\leq k_1\leq n} \left(\sum_{s_1\geq s} \frac{X_{(k_1,s_1)}}{n} + \sum_{s_1\geq s+1} \frac{X_{(k_1,s_1)}}{n} \right).\end{aligned}$$

For ease of notation we define $g^{(n)}_s := \sum_{1\leq k_1\leq n} \sum_{s_1=s}^{b} \frac{X^{(n)}_{(k_1,s_1)}}{n}$ and so we get full transition rates

$$X^{(n)} \to X^{(n)} - e^{(n)}_{(k,s)} + e^{(n)}_{(k,s-1)} \quad \text{at rate } \mu_k X^{(n)}_{(k,s)} \mathbb{1}_{\{s>0\}} \tag{3.3}$$

$$X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)} \quad \text{at rate } \lambda X^{(n)}_{(k,s)} (g^{(n)}_s + g^{(n)}_{s+1}). \tag{3.4}$$

Notice that $g^{(n)}_s$ gives the proportion of n servers that have queue length s or greater. This is a similarity to the classical mean field analysis in Section 2.3.1 and has a similar purpose since it represents the probability of choosing a second server with a longer queue.

3.2 Drift and mean field

In order to find the mean field we first must find the drift. In this case it is a matrix valued function with a matrix output and for each (k, s) (Using definition 2) for some state $x^{(n)} \in \text{conv}(\mathcal{X}^{(n)})$ it is given by

$$f^{(n)}_{(k,s)}(x^{(n)}) = \mu_k x^{(n)}_{(k,s+1)} - \lambda x^{(n)}_{(k,s)} (g^{(n)}_s + g^{(n)}_{s+1}) - \left(\mu_k x^{(n)}_{(k,s)} - \lambda x^{(n)}_{(k,s-1)} (g^{(n)}_{s-1} + g^{(n)}_s) \right) \mathbb{1}_{\{s>0\}}. \tag{3.5}$$

We remember that the drift shows the expected change in the system in some infinitesimal time δ. Since we have a Markov process, this means this is equal to the sum of changes in state multiplied by transition rate. The binary values in the form $X^{(n)}_{(k,s)}$ act as indicators that server k has queue length s. If server k is of queue length $s + 1$ then we have departures at rate μ_k which contributes an expected change to length s of μ_k. Likewise departures from queue length of $s > 0$ cause an expected change of $-\mu_k$. If the queue length is s, then an arrival causes an expected change of $-\lambda(g^{(n)}_s + g^{(n)}_{s+1})$ and similarly an arrival to a queue of length $s - 1$ causes an expected change of $\lambda(g^{(n)}_{s-1} + g^{(n)}_s)$.

We use the drift to construct the mean field approximation and we find that for any initial condition $x \in \mathcal{X}^{(n)}$ we have $\phi : \mathcal{X}^{(n)} \times \mathbb{R}^+ \to \mathbb{R}^{n\times|S|}$ that satisfies

$$\frac{d}{dt}\phi^{(n)}(x, t) = f^{(n)}(\phi^{(n)}(x, t)) \tag{3.6}$$

gives the mean field. However, we cannot solve this in general and so we turn to numerical methods discussed in Section 3.4.

Consider a sequence of systems $X = \left(X^{(d(n))}\right)_{n\in\mathbb{N}}$ with $d(n)$ the system size of each element. In order to apply Theorem 2 to bound the error term along the sequence X, we must check the bounds on the partial transition rates for each system $X^{(d(n))}$. Since we choose finite processing rates independent of $d(n)$, the partial rates (3.1) are $O(1)$ bounded. Then for the partial arrival rates (3.2), we have that $X^{(d(n))}_{(k,s)} \leq 1$ and so $(2\lambda\mathbb{1}_{\{s_1>s\}} + \lambda\mathbb{1}_{\{s_1=s\}})/d(n)$ is clearly $O(1/d(n))$ bounded. Therefore we can apply Theorem 2 as long as we have finite processing rates.

3.3 Fixed proportions of service rates

We would like to consider sequences of growing systems and analyse the behaviour of their mean fields. An arising issue is how to define the growth of a system when we have heterogeneous processing rates. A case where the growth is well defined is when each system in the sequence has the same proportion of different speed servers. An example of such a system can be seen in Figure 3.1. Before we consider sequences of systems, we first look at just scaling one fixed system which we more precisely define in the following definition.

Definition 4 (Proportionally Scaled System)**.** Consider a system $X^{(N)}$ of N servers described by the model in Section 3.1 with heterogeneous service rates μ_k, arrival rate λ and JSQ(2) on the complete graph with replacement. Suppose we group each server with others of the same processing rate also known as type. Let m be the number of groups and the set of servers in each group be given by $V_l^{(N)}$ each with processing rate μ^l and the number of servers in group l be given by $\left|V_l^N\right| = \alpha_l N$. Then define the degree a Proportionally Scaled System of $X^{(N)}$ to be a second system $\tilde{X}^{(aN)}$ which has aN servers with the same groups $\tilde{V}_l^{(aN)}$ of rate μ^l but now of size $|\tilde{V}_l^{(aN)}| = \alpha_l aN$. Unless stated otherwise, also assume the initial condition is scaled such that any two servers $k \in V_l^{(N)}, k' \in V_l^{(aN)}$ share the same initial condition

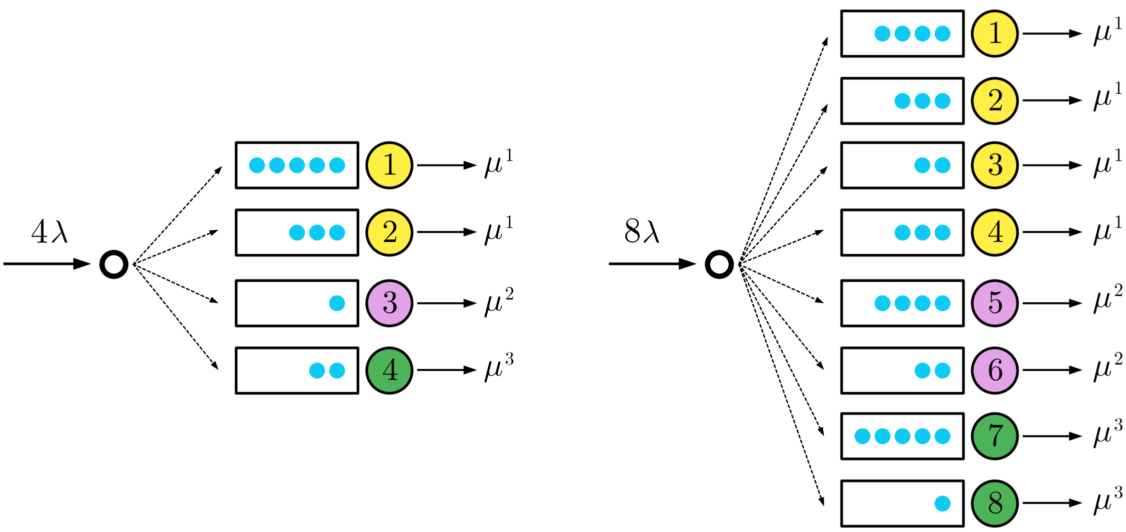

Figure 3.1: Initial system on the left with $m = 3$ groups and proportionally scaled system with $a = 2$ on the right.

Remark 1. In Definition 8 we use the factor a to simplify the growth. Taking integer multiples of the first system size means that we ensure that $\alpha_l aN$ is an integer for all l. We can then view this growth in two ways, either adding $(a-1)N$ servers to the system or resampling a new system of size aN, both have the same result because the service rates and structure are deterministic.

3.3.1 Finite systems

The new mean field considers the behaviour of each of the n servers independently and so perhaps there is a natural suggestion that the approximation depends on n which is also seen in the notation e.g. $\phi^{(n)}$ from Theorem 2. We suggest this is not the case when there are fixed proportions of servers with given rates. In Theorem 3 we explain that the proportionally scaled systems are independent of n.

Theorem 3. *Consider a system $X^{(N)}$ and proportionally scaled system $\tilde{X}^{(aN)}$ for any $a \in \mathbb{N}$ (Definition 4) and two arbitrary servers of the same type from both systems, $k_1 \in \tilde{V}_l^{(aN)}$, $k \in V_l^{(N)}$. If initial conditions are equal for all servers of the same type (defined as x^l for type l), then $\tilde{\phi}^{(aN)}_{(k_1,s)}(\tilde{x},t) = \phi^{(N)}_{(k,s)}(x,t)$. Furthermore, $\phi^{(N)}_{(k,s)}(x,t) = \psi_{(l,s)}(z,t)$ where ψ is the unique solution of the system of ODEs*

$$\frac{d\psi_{(l,s)}}{dt} = \mu^l(\psi_{(l,s+1)} - \psi_{(l,s)}\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_{(l,s)}(g_s + g_{s+1}) - \psi_{(l,s-1)}(g_{s-1} + g_s)\mathbb{1}_{\{s>0\}}\right) \tag{3.7}$$

with initial condition $z = (x^1, ..., x^m)$ and $g_s = \sum_{s' \geq s} \alpha_l \psi_{(l,s)}$.

The full proof of Theorem 3 can be found in Appendix A.2.1 and follows proof plan A.2. However, we sketch the proof here. The idea is that we notice servers of the same type, starting from the same initial condition have the same evolution and so have equal mean field for all time. We use this 'exchangeability' between servers to obtain a more classical mean field analysis using proportions of queue lengths of exchangeable servers. From here we eliminate the dependency on n by collapsing the system of n equations to m, one for each server type. The intuition for this result follows from the intuition for the classical mean field. The form with proportions of servers of the same type could be constructed from the perspective of the classical mean field, see Corollary 3. As a corollary, it is also clear that if we took $m = 1$ in (A.3), in other words the homogeneous case, and sum over s then we obtain the classical mean field ODE (2.4) (only with proportions of servers of length s rather than greater or equal to s).

We now make some remarks and corollaries with regard to Theorem 3 in order to explain some of its implications.

Remark 2 (Mean field interpretation)**.** Although our goal is to consider sequences of systems, Definition 4 and Theorem 3 do not explicitly refer to sequences. We interpret the mean field in a particle sense, allowing us to find the mean field of these fixed models.

Remark 3 (Independence of n)**.** Since the new mean field approximation has a state space dependent on n it is not easy to see how to compare the mean field between two differing sized systems. In this case we grouped together servers with identical behaviour and thus we can compare the state space of groups between different sized systems. Proportional scaling then means that the structure of the heterogeneity in the system is preserved as it grows in size because the relative sizes of all the groups stay the same. Theorem 3 suggests exactly this by saying that any two servers with the same processing rate from a system and a proportionally scaled system (Definition 4) have the same mean field.

Remark 4. Finally we note that Theorem 3 does not hold if we use the JSQ(2) policy without replacement. The neighbourhood of a given server is the set of all other servers except itself. Therefore if we choose a server of type l then the proportion of type l servers in its neighbourhood is $(\alpha_l n - 1)/(n-1)$ and the proportion of type $j \neq l$ is $\alpha_j n/(n-1)$. These proportions are not constant and depends on n. This means when we sample the second server, the probability to choose a given server type depends on n. In effect, we no longer have fixed proportions. Although, if we take the limit of these proportions we see $(\alpha_l n - 1)/(n-1) \to \alpha_l$ and $\alpha_j n/(n-1) \to \alpha_j$ as $n \to \infty$ and so in the limiting system we have fixed proportions. Another way of looking at this is to consider that the probability to sample the same server twice tends to zero as $n \to \infty$ and so replacement has no effect. In Section 5.3 we will consider a type of scaling on graphs for JSQ(2) without replacement such that we obtain a similar result to Theorem 3 whereby the mean field is independent of n.

3.3.2 Scaling results

In Section 3.3.1 we considered a particle like interpretation of the new mean field by looking at finite systems. We now wish to make limiting statements about the behaviour of proportionally scaled systems as $n \to \infty$. The aim of this section is to connect the particle and scaling-like interpretations of the mean field posed in Section 2.3.3 for proportional scaling.

Using Definition 4 we construct sequences of proportionally scaled systems $\left(X^{(d(a)N)}\right)_{a\in\mathbb{N}}$ where $(d(a))_{a\in\mathbb{N}}$ with $d(a) \in \mathbb{N}$ such that $d(a) \to \infty$ as $a \to \infty$. Then the a-th element of the sequence is the proportionally scaled system of $X^{(N)}$ of degree $d(a)$. However, it is difficult to derive a meaning from the limit as $a \to \infty$ because the state space of the a-th model is of dimension $d(a) \times b$ and so growing with a. In order to make limiting statements we must use a quantity that is comparable for systems of different size. In the proof of Theorem 3 we used $\psi_{(l,s)}$ which represents the proportion of servers of type l that have queue length s in the mean field. We formalise this notion in the following definition

Definition 5 (Occupancy state)**.** Consider a system $X^{(n)}$ with m groups of servers $V_l^{(n)}$ of size $\alpha_l n$ that share processing rate μ^l as described in Definition 4. Then let $x_{(l,s)}^{(n)}(t)$ be the proportion of servers of type l that have queue length s at time t which is defined as

$$x_{(l,s)}^{(n)}(t) := \frac{1}{\alpha_l n} \sum_{k \in V_l^{(n)}} X_{(k,s)}^{(n)}(t).$$

Then we define the full occupation state as such $x^{(n)}(t) := (x_{(l,s)}^{(n)}(t))_{l=1,..,m,\ s\in\mathcal{S}}$.

The quantity $x_{(l,s)}^{(n)}(t) \in [0,1]$ for all $t \in \mathbb{R}^+$ and for any $n \in \mathbb{N}$ because it is scaled by dividing through by the group size $|V_l^{(n)}| = \alpha_l n$. Therefore this quantity is well defined for all n and we can look at its limit as $n \to \infty$, although it remains to be shown that the limit exists. The limit in distribution of the occupancy state is defined as $x_{(l,s)}$ which is known as the fluid limit. Finding this directly is a more involved problem and discussed in the scaling limits approach in Section 2.3.1. Note that when we drop the dependence on t the occupancy state is a function $x_{(l,s)}^{(n)} : \mathbb{R}^+ \to [0, 1]$. We look at a weaker quantity than the fluid limit, the limit of the expected occupancy state (LEOS) in the following result.

Theorem 4 (LEOS)**.** *Consider a sequence of proportionally scaled systems $(X^{(d(a)N)})_{a\in\mathbb{N}}$, where each system has occupancy state given by $x_{(l,s)}^{(d(a)N)}$ for servers of type l with queue length s (Definition 5). If each server of the same type l for any system in the sequence have the same initial condition x_l, then the expectation of the occupancy state converges as such*

$$\mathbb{E}\left[x_{(l,s)}^{(d(a)N)}(t)\right] \to \psi_{(l,s)}(t, z), \quad as \;\; a \to \infty$$

where ψ is the unique solution of the system of ODEs

$$\frac{d\psi_{(l,s)}}{dt} = \mu^l (\psi_{(l,s+1)} - \psi_{(l,s)}\mathbb{1}_{\{s>0\}}) - \lambda \left(\psi_{(l,s)}(g_s + g_{s+1}) - \psi_{(l,s-1)}(g_{s-1} + g_s)\mathbb{1}_{\{s>0\}}\right) \quad (3.8)$$

with initial condition $z = (z_1, ..., z_m)$. We call $\psi(t, z)$ the LEOS of $(X^{(d(a)N)})_{a\in\mathbb{N}}$.

A full proof can be found in Appendix A.3.1. The idea is to use linearity of the expectation to analyse the expectation of $X_{(k,s)}^{(n)}$ for $k \in V_l^{(n)}$ through Theorem 2. This links the expectation to a sum of the mean field of servers of the same type. From Theorem 3 and its proof we know that servers of the same type have the same mean field which uniquely satisfies (3.8). We now give some remarks and implications of this result.

Remark 5 (LEOS)**.** We make it very clear that the LEOS is not a fluid limit. The fluid limit is the limit in distribution of the occupancy process, whereas we have the limit of the expected occupancy state. We are left with the expectation as a product of using the mean field in our analysis. The mean field approach is weaker than the scaling limits approach as given in the classical case in Section 2.3.1. Since the mean field considers the expected change in the system, it is somewhat natural that we consider the expectation of the fluid limit here. In order to say the limit of the expectation is equal to the expectation of the fluid limit we must be able to take the limit inside the expectation. This would require knowing that the fluid limit exists and for example using the dominated convergence theorem. A priori there is no reason why the fluid limit should exist in general. Should the fluid limit exist and we take the limit inside the expectation then the expected fluid limit is generally equal to the fluid limit since it should be deterministic in this case.

Remark 6 (Interpretation of mean field)**.** Theorem 4 states that the LEOS is the unique solution of the system of ODEs (3.8) which is the same system found for finite systems in Theorem 3. This means that the particle interpretation of the mean field for any finite system is equal to a scaling-like interpretation with the LEOS. This helps us to understand that the alternative interpretations for the classical mean field posed in Section 2.3.1 also exist for the new mean field in this case of proportional scaling. This is an idea explored in Section 2.3.3.

We formalise the notion in Remark 6 with the following corollary.

Corollary 1 (Finite mean field as a LEOS)**.** *Take any system $X^{(N)}$ as described in Section 3.1. Group the servers with the same processing rate into m groups $V_l^{(N)}, l = 1, ..., m$. Then the new mean field of server $k \in V_l^{(N)}$ is given by the LEOS of the sequence of proportionally scaled systems $(X^{(aN)})_{a\in\mathbb{N}}$ (Definition 4).*

Proof. The new mean field of the finite system $X^{(N)}$ for a server $k \in V_l^{(N)}$ and queue length s is given by $\phi^{(N)}_{(k,s)} = \psi_{(l,s)}$, the unique solution of ODE (3.7) by Theorem 3. The LEOS of the sequence of proportionally scaled systems $(X^{(aN)})_{a\in\mathbb{N}}$ is given by the unique solution of the same system of ODE as per Theorem 4 and since the solutions are unique they must be equal. □

Remark 7. The result of Corollary 1 is subtle. It says that if we take any heterogeneous system on the complete graph as explained in Section 3.1, then its particle-like new mean field can also be interpreted as the LEOS of a proportionally scaled system starting from the initial system. Another way to see it is that for every finite system, there exists an 'implicit' sequence of proportionally scaled systems such that the finite mean field 'equals' the LEOS of the implicit sequence. By equals we must be careful because the state spaces could be different sizes if we group servers of the same size. This issue can be avoided by saying in the initial system we have N groups, one for each server. Therefore state spaces will be the same $N \times b$ size and so we can say they are equal.

3.3.3 Power-of-d

In this work we focus on the power-of-two or JSQ(2) policy in order to limit the number of interacting objects because the new mean field analysis becomes more cumbersome with more interacting objects. In [3] they provide results for d interacting objects but for the most part restrict themselves to two interacting objects. A relevant question is whether our work is particular to power-of-two or does it extend to power-of-d. In this Section we briefly discuss the mean field of fixed proportion systems using JSQ(d) with replacement.

We take a classical population process approach to finding the mean field. Let $q_{(l,s)}$ be the proportion of type l servers that have queue length s or greater. Then we find $q_{(l,s)}$ satisfies

$$\frac{dq_{(l,s)}}{dt} = \lambda P_{(l,s-1)} - \mu^l \left(q_{(l,s)} - q_{(l,s+1)}\right) \tag{3.9}$$

$$P_{(l,s)} = \sum_{i=1}^{d}\sum_{j=1}^{d-i} \frac{i}{i+j} B(d,j,i) p_{s+1}^{d-i-j} (q_{(l,s)} - q_{(l,s+1)})^i ((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))^j \tag{3.10}$$

$$p_s = \sum_{l=1}^{m} q_{(l,s)}, \qquad B(d,j,i) = \frac{d!}{i!(j-i)!(d-j)!} \tag{3.11}$$

The derivation is very similar to the population process in Section 2.3.1 and can be found in full in Appendix B.1. We now discuss a few corollaries to look specific cases of this scheme, first verifying that this aligns with the literature in [1].

Corollary 2. *In the homogeneous case, the system of ODEs* (3.9) *are equal to the system following system*

$$\frac{dq_s}{dt} = \lambda(q_s^d - q_{s+1}^d) - \mu\left(q_s - q_{s+1}\right). \tag{3.12}$$

Proof. In the homogeneous case we have $m = 1$ and so $p_s = q_{(1,s)} := q_s$ as well as $\mu^l = \mu$. Looking at $P_{(l,s)}$ we see that the sum over j is zero unless $j = 0$. We also see that

$B(d,i,0) = \binom{d}{i}$ and so

$$P_{(l,s)} = \sum_{i=1}^{d} \binom{d}{i} q_{s+1}^{d-i}(q_s - q_{s+1})^i = \left(\sum_{i=0}^{d} \binom{d}{i} q_{s+1}^{d-i}(q_s - q_{s+1})^i \right) - q_{s+1}^d.$$

We recognise the summation term as the binomial expansion of $(q_{s+1} - (q_s - q_{s+1}))^d = q_s^d$. Thus $P_{(l,s)} = (q_s^d - q_{s+1}^d)$ and substitution into (3.9) gives (3.12). □

Corollary 3. *For the $d = 2$ power-of-two scheme, the system of ODEs (3.9) can be obtained by the system (3.7) with $q_{(l,s)} = \sum_{s' \geq s} \psi_{(l,s')}$.*

Proof. For $d = 2$ we find

$$\begin{aligned} P_{(l,s)} &= 2p_{s+1}(q_{(l,s)} - q_{(l,s+1)}) + 2 \cdot \frac{1}{2}((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))((q_{(l,s)} - q_{(l,s+1)})) \\ &\quad + 1 \cdot 1 \cdot (q_{(l,s)} - q_{(l,s+1)})^2 \\ &= 2p_{s+1}(q_{(l,s)} - q_{(l,s+1)}) + (p_s - p_{s+1})((q_{(l,s)} - q_{(l,s+1)})) \\ &= (p_{s+1} + p_s)(q_{(l,s)} - q_{(l,s+1)}). \end{aligned}$$

Substitution into (3.9) yields

$$\frac{dq_{(l,s)}}{dt} = \lambda(p_s + p_{s-1})(q_{(l,s-1)} - q_{(l,s)}) - \mu^l \left(q_{(l,s)} - q_{(l,s+1)} \right). \tag{3.13}$$

This considers the proportions of servers with queue lengths greater than or equal to s whereas the system in (3.7) only considers queue lengths equal to s. Therefore we sum (3.7) over s. We see that there are many telescoping sums and so we obtain

$$\sum_{s' \geq s} \frac{d\psi_{(l,s)}}{dt} = -\mu^l \psi_{(l,s)} - \lambda \psi_{(l,s-1)}(g_{s-1} + g_s).$$

Notice that $q_{(l,s)} = \sum_{s' \geq s} \psi_{(l,s')}$ and therefore $\psi_{(l,s)} = q_{(l,s)} - q_{(l,s+s)}$ and furthermore $g_s = q_s$ which upon substitution yields exactly (3.13). □

The main implication of Corollary 3 is that our earlier system of ODEs (2.4) (which describes the mean field of proportionally scaled systems from a new mean field perspective) can also be found with a population process method. In order to construct (3.9) we say the ODE is exact as $n \to \infty$ and so this is akin to Remark 7. From this we see how the finite new mean field is equal to the population process as n goes to infinity. Now we have shown that the new mean field can be interpreted in the same three ways as the classical mean field in Section 2.3.1 for proportionally scaled systems.

3.4 Numerical Results

Although the main focus of this work is theoretical, we provide some brief numerical results to illustrate our theoretical results with examples. In [17] Allmeier and Gast provide a method and tool for the computation of mean fields and refined mean fields. This tool is generalised for any model that can be described as in Section 2.3.2 and not only for load balancing problems. We use our own, less generalised implementation, to produce numerical results. In all sections on numerical results we consider figures of the average queue length over time. This does not directly show the mean field for each server type which is a probability, however it demonstrates the key trends in a more concise manner.

In the following example we demonstrate proportional scaling. We have the fully connected graph using JSQ(2) with replacement, arrival rate $\lambda = 1$ and four groups of unique service rates $\mu^1, \mu^2, \mu^3, \mu^4 = 2, 0.5, 1.4, 1$ with corresponding proportions $\alpha_1, \alpha_2, \alpha_3, \alpha_4 = 1/5, 1/5, 1/5, 2/5$. We both simulate the system and compute the mean field for $n = 10, 20, 30, 40$ and for the computation of the mean field we use a buffer of size 12 and no buffer for the simulation. With this buffer size the drop rate is negligible for all systems and thus a good approximation of the simulation with no buffer. It may appear that the buffer should grow with n, but we know that the mean field for servers of the same type are equal across all systems and so if its sufficient for 10, then it is also sufficient for 40. The simulation is performed 2500 times and the mean result is given with 95% Wald confidence interval.

We present the average queue length over time in Figure 3.2. Immediately we notice that the average queue length for the mean field is the same for each n. This is certainly expected based on Theorem 3. In the proportionally scaled system the mean field of servers of the same type is the same and the number of servers of each type scale linearly with n thus the mean is unchanged.

From Theorem 2 we know the error term between the stationary distribution and mean field is $O(1/n)$ and therefore the error between the average queue length of the simulation and mean field should also be $O(1/n)$. This is visible in Figure 3.2 as we see the average queue length of the simulation tends to that of the mean field. This also appears to happen in a $O(1/n)$ fashion with diminishing returns as n increases. As discussed at the end of Section 2.3.2, the constant associated with the error term depends on the specific system. Therefore for small n such as in Figure 3.2 the errors may fluctuate and not appear to follow an $O(1/n)$ trend. However, this is not observed in this example, likely because the proportionally scaled systems are sufficiently similar (preserve heterogeneity) such that the associated constants for the error terms are similar between systems.

Next we notice that the 95% confidence interval becomes much smaller as n grows despite each having the same number of runs. This happens because the variance in the simulation is reduced. As $n \to \infty$ the system should behave deterministically and so the variance in the simulation tends to zero as n grows large.

This example was inspired by an example in [3, Section 5.2.2]. In their example they used a mixture of fixed proportions of service rates with one fifth having rate 2 and another fifth with rate 0.5 with the remaining three fifths of servers having a rate uniformly sampled between 1 and 1.4. They then sampled four systems of size 10, 20, 30, 40 on which they performed the simulation and computed the mean field. Our initial intention was to replicate their example but this was difficult since their example is essentially one random sample.

For the systems they sampled, they found that the mean field was not constant and that the average queue length of the mean field decreased as n grew. As their systems grow, there are more samples of rates and so the distribution of the randomly sampled rates should become more even. We expect this to benefit the system since the negative impact of a slow

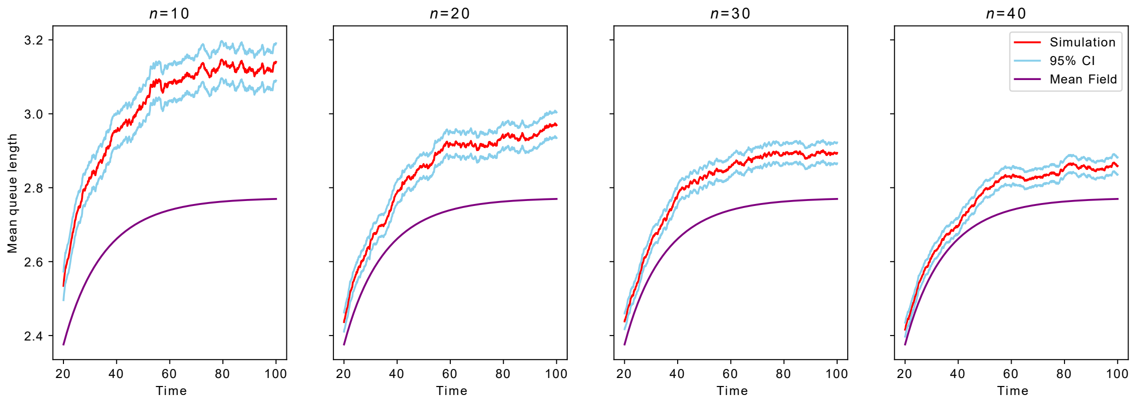

Figure 3.2: Average queue length of simulation vs mean field approximation for a proportionally scaled system.

server should be proportionally greater than the positive impact of a fast server. Therefore as n increases, the probability to sample a better system increases and so this heuristically explains why the average queue length of the mean field decreases as n increases in their example. Despite the similarities in rates between both examples, the random examples are not proportionally scaled and so we have very different systems each time resulting in a different mean field. We could proportionally scale a randomly sampled system which would then have the same mean field, but then this is no longer representative of the random system.

Chapter 4

Servers on graph structures

In this section we connect the new mean field method with the classical mean field methods on graphs with homogeneous processing rates found in [12]. Using the homogeneous case allows us to focus on the impact of the graph structure. We continue two of the main themes from Chapter 3 by again looking at when the mean field is 'independent' of the system size n and trying to find the LEOS for sequences of growing systems.

4.1 Model

We use the model outlined in Section 2.2. For a system of n servers we have Poisson arrivals at rate λn and exponential processing times with homogeneous processing rate μ. We apply JSQ(2) on graphs to some arbitrary graph $G_n = (V_n, E_n)$. An example 'arbitrary' graph can be seen on the right hand side of Figure 2.1. In this section we will find that the mean field is the same for JSQ(2) with and without replacement, however the model without replacement is leading. The stability conditions for this model could be derived from results in [20], for the purpose of our work we assume the system is stable.

We use notation from Section 2.3.2 and so the state space is given by the $n \times b$ matrix of binary objects $X^{(n)}$. We first describe the transition rates of the model. For any two servers $k, k_1 \in V_n$ and queue lengths $s, s_1 \in S$ of servers k, k_1 respectively, we have partial transition rates

$$X^{(n)} \to X^{(n)} - e^{(n)}_{(k,s)} + e^{(n)}_{(k,s-1)} \quad \text{at rate } \mu X^{(n)}_{(k,s)} \mathbb{1}_{\{s>0\}} \tag{4.1}$$

$$X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)} \quad \text{at rate } (\lambda \mathbb{1}_{\{s_1>s\}} + \lambda/2 \mathbb{1}_{\{s_1=s\}})\alpha \mathbb{1}_{\left\{k_1 \in \mathcal{N}^{(n)}_k\right\}}, \tag{4.2}$$

where α is given by

$$\alpha = \frac{X_{(k,s)}}{n} \frac{X_{(k_1,s_1)}}{\left|\mathcal{N}^{(n)}_k\right|} + \frac{X_{(k_1,s_1)}}{n} \frac{X_{(k,s)}}{\left|\mathcal{N}^{(n)}_{k_1}\right|}.$$

Much like earlier analysis we find it useful to define a $g^{(n)}_{(k,s)}$,

$$g^{(n)}_{(k,s)} := \frac{1}{2} \sum_{k_1 \in \mathcal{N}_k} \sum_{s_1=s}^{b} X_{(k_1,s_1)} \left(\frac{1}{\left|\mathcal{N}^{(n)}_k\right|} + \frac{1}{\left|\mathcal{N}^{(n)}_{k_1}\right|} \right).$$

By summing over all (k_1, s_1) we obtain the full transition rates. They are in the same form as Section 3.1 transition rates (3.1), (3.2) but with this different $g^{(n)}_{(k,s)}$ as such

$$X^{(n)} \to X^{(n)} - e^{(n)}_{(k,s)} + e^{(n)}_{(k,s-1)} \quad \text{at rate } \mu X^{(n)}_{(k,s)} \mathbb{1}_{\{s>0\}} \tag{4.3}$$

$$X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)} \quad \text{at rate } \lambda X^{(n)}_{(k,s)} (g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}). \tag{4.4}$$

We note that when we have the fully connected graph, $|\mathcal{N}^{(n)}_k| = n$ for all k and then clearly $g^{(n)}_{(k,s)} = g^{(n)}_s$ from Section 3.1. This then gives us the homogeneous case of the model in Section 3.1. The key difference is that $g^{(n)}_{(k,s)}$ now depends on k since we only consider its neighbourhood.

4.2 Drift and Mean field

The derivation for the drift follows the same logic as in Section 3.2 and due to the almost identical transition rates, we obtain an almost identical drift to (3.5). The differences are that we have our more generalised definition of $g^{(n)}_{(k,s)}$ and homogeneous processing rate μ. We give the drift below for a state $x^{(n)} \in \text{conv}(\mathcal{X}^{(n)})$.

$$f^{(n)}_{(k,s)}(x^{(n)}) = \mu x^{(n)}_{(k,s+1)} - \lambda x^{(n)}_{(k,s)} (g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}) - \left(\mu x^{(n)}_{(k,s)} - \lambda x^{(n)}_{(k,s-1)} (g^{(n)}_{(k,s-1)} + g^{(n)}_{(k,s)})\right) \mathbb{1}_{\{s>0\}}. \tag{4.5}$$

In order to apply Theorem 2 we must satisfy the condition on the bounds of the partial transition rates (4.1), (4.2) for a sequence of systems $(X^{(d(n))})_{n\in\mathbb{N}}$. We see that μ is clearly $O(1)$ bounded and thus so is (4.1). However the partial arrival rates (4.2) are not necessarily $O(1/d(n))$ bounded. We see that

$$(\lambda \mathbb{1}_{\{s_1>s\}} + \lambda/2\mathbb{1}_{\{s_1=s\}}) \left(\frac{1}{\left|\mathcal{N}^{(d(n))}_k\right|} + \frac{1}{\left|\mathcal{N}^{(d(n))}_{k_1}\right|} \right), \qquad \text{for all } k, k_1 \in V_{d(n)}$$

is only $O(1/d(n))$ bounded if $\left|\mathcal{N}^{(d(n))}_k\right| = O(d(n))$ for all $k \in V_{d(n)}$. This means that every neighbourhood in the sequence of systems must grow with $d(n)$, the number of servers in the n-th system.

4.3 Growth of neighbourhoods

As mentioned at the end of Section 4.2 if we take a sequence of systems then Theorem 2 only applies if the size of the neighbourhoods of all servers grow with $O(n)$ where n is the number of servers in the system. Much like in Section 3.3.2 we would like to analyse the behaviour of a sequence of system as $n \to \infty$. In order to do this we make use of Theorem 2 because as $n \to \infty$ the $O(1/n)$ error term vanishes and so we can construct limiting results. However, we only need an error term that vanishes as $n \to \infty$ and so we consider slower growing neighbourhoods with the expectation that this yields a larger but still vanishing error term.

Consider a sequence of systems according to Section 4.1, $X = \left(X^{(n(a))}\right)_{a\in\mathbb{N}}$ and so the a-th system is of size $n(a)$ with graph structure $G_{n(a)}$, where $(n(a))_{a\in\mathbb{N}}$ an integer sequence with

$n(a) \to \infty$ as $a \to \infty$. Let $\phi^{(n(a))}$ denote the mean field for the a-th system in the sequence, then we obtain the following proposition

Proposition 1. *Consider the model as described in Section 2.2 and the above sequence. Let $\phi^{(n(a))}(x,t)$ be the solution of (2.11) with drift $f^{(n(a))}$ and initial condition $x \in conv(\mathcal{X}^{(n(a))})$ for the $a-th$ element of the sequence. If the processing rates μ_k are $O(1)$ bounded and the arrival rates to each server pair are $O(1/n(a)^p)$ bounded $p \in (0,1]$ then we have for $(k,s) \in \{1,...,n\} \times S$ and $t < \infty$*

$$\mathbb{P}\left(S_k^{(n(a))}(t) = s\right) = \mathbb{E}\left[X_{(k,s)}^{(n(a))}(t)\right] = \phi_{(k,s)}^{(n(a))}(x,t) + O(1/n(a)^p),$$

where $S_k^{(n(a))}(t)$ is the random variable which gives the queue length of server k at time t of the $a-th$ system in the sequence.

A full proof of Proposition 1 requires a complete replication of the proof of [3, Theorem 2.2] which is a long and technical proof. We therefore do not provide a full proof but rather a guide of how the proof for [3, Theorem 2.2] can be adapted for our specific case to yield this result.

Proof plan. Consider the proof from [3, Section 6.3] the only difference is in bounding the error term which requires looking at [3, Lemma D.2]. This is described as the most technical lemma in [3] and we describe the steps to prove this as such:

- First consider the order of our transition rates. For simpler notation let $n = n(a)$ be the system size. In their notation, our transition rates for just one interacting object $r_{k,s\to s'}$ are O(1) bounded (4.1). Our transition rates for two interacting objects $r_{k,k',(s,s')\to(s_1,s_1')}$ must be $O(n^{1-p})$ if the partial arrival rates (4.2) are $O(1/n^p)$ bounded because the r is given with a factor of $1/n$ in their notation.
- The error term is a sum over all binary states of the product of the expected change of the covariance in the system and the second derivative of mean field with respect to the binary states.
- First we find the expected change of the covariance of the stochastic system. This is split into three cases. The covariance of objects with the same server and queue length $(k,s),(k,s)$ is $O(1)$, likewise for same server different queue lengths $(k,s),(k,s')$. However, for two different servers $(k,s) \neq (k',s')$ we find this is $O(1/n^p)$. This is because the only interactions between two objects is a partial arrival which is $O(1/n^p)$ and so we lose the $O(1)$ term of departures. Therefore we find the covariance $Q_{(k,s)(k_1,s_1)} = O(c_{k_1}^k)$ where $c_{k_1}^k = 1$ if $k = k_1$ and $1/n^p$ otherwise.
- We now consider bounding the second derivatives of the mean field $\phi_{(k,s)}^{(n)}$ with respect to the indicator states. This is the most technical part of the proof. They look at the derivative of the mean field with respect to time which gives the drift. Notice that n dictates the size of the state space and thus the summation limit too, however as a term it only appears as a constant in the partial transition rates. We have the same size n, our different factor $O(n^p)$ is only found in the transition rates as a constant and so persists through the calculations. Essentially, none of the steps to bound the derivative in [3, Lemma D.1] involve n. This means we end up replacing the $c_{k_1}^k$ at the end with

our redefined version that has $O(1/n^p)$ rather than $O(1/n)$. We then use this bound of the first derivative to obtain a bound on the second derivative. This follows similar reasoning to the first derivative that n is only found as a constant. Thus we obtain the bound for the second derivative with respect to servers k_1, k_2 as such $O((c^k_{k_1} + c^k_{k_2})/n^p)$.

- Returning to [3, Lemma D.2] we find the product and sum over the covariance and second derivatives is the same with out new definition of $c^k_{k_1}$ for which it can be shown gives a final bound of the error term R as $O(1/n^p)$.

□

The main implication of Proposition 1 is that we can weaken the condition on the growth of neighbourhood sizes in the end of Section 4.2. In this case we can have sequences of models with neighbourhoods that grow with order $O(n(a)^p),\ p \in (0, 1]$ following the same reasoning from the aforementioned section. This means the neighbourhoods can grow much more slowly, allowing us to consider a wider range of sequences. Note that the sequence of system sizes $(n(a))_{a\in\mathbb{N}}$ is used to define a more general sequence of models, but that we see the neighbourhoods must grow with $n(a)$ and thus the size of the system. If $n(a)$ grows more rapidly then the error term will also decrease more rapidly, however the neighbourhood sizes will also grow more rapidly. This is a means to say that the specific sequence chosen has no affect on the conditions of the relative neighbourhood sizes.

Briefly we explain how this extension of Theorem 2 has an intuitive explanation. The system described by the mean field behaves as though all of the servers evolve independently [3]. When the neighbourhood sizes grow to infinity, the change in state of any finite number of neighbours has no effect on the evolution of the principal server of the neighbourhood. Thus as long as the neighbourhood sizes tend to infinity, the change of one server is washed out by the sheer size of infinity. Therefore the neighbourhoods can grow as slowly as they want as long as they grow to infinity.

Remark 8. We also believe that the condition of $O(n(a)^p), p \in (0, 1]$ growing neighbourhoods is still too strict. For example it does not allow for logarithmic growth. If we consider the proof plan there is no reason why we cannot allow for any neighbourhood growth as long as they tend to infinity. The aforementioned intuitive explanation holds too. To use Proposition 1 precisely we should require $O(n(a)^p), p \in (0, 1]$ neighbourhoods but for brevity we will use the more relaxed condition suggested here that the neighbourhood sizes tend to infinity. In all cases that we use this proposition with the relaxed condition, it would not fundamentally change the result to enforce $O(n(a)^p), p \in (0, 1]$ neighbourhoods.

4.4 Regular graphs

In this section we study systems with regular graph structures and homogeneous processing rates. We use this as a first step to understand the mean field on general graph structures. In regular graphs every node has the same degree and so we have a nonincreasing set of parameters. We will see that due to the homogeneous service rates and equal size neighbourhoods, every neighbourhood is identical and so we obtain an exchangeability result. Regular graphs mean that for a given model all of the interaction neighbourhoods are of the same size i.e. $|\mathcal{N}^{(n)}_k| = d$ for all $k \in V_n$ where d is called the degree. Two examples of regular graphs can be seen in Figures 4.1 and 4.2. This section follows a similar structure to Section 3.3. We first look at the new mean field of finite systems and then use this to construct limiting results.

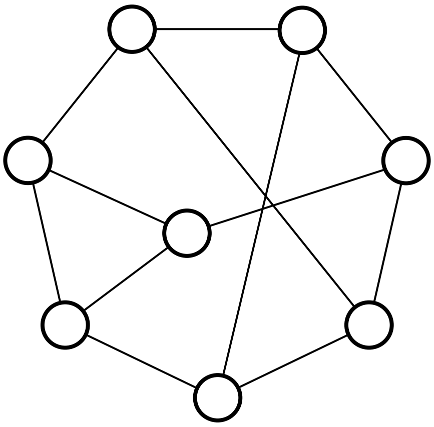

Figure 4.1: Regular graph with $n = 7$ and $d = 3$.

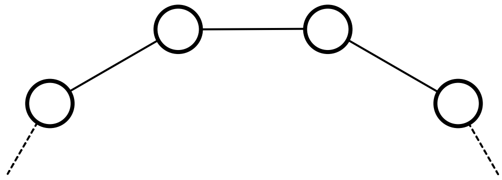

Figure 4.2: Regular graph of size n with $d = 2$ known as a ring.

4.4.1 Finite systems

The new mean field considers the behaviour of every server individually and so there is perhaps a natural suggestion that this captures the structure between servers encoded in the graph. In [11] Gast discusses that the classical mean field does not provide a good approximation on graph structures. In the following theorem we show that the new mean field on any regular graph satisfies the same system of ODEs and so the new mean field also appears not to capture the structure of the graphs.

Theorem 5. *Consider a system $X^{(n)}$ as described in Section 4.1 with a regular graph structure $G_n = (V_n, E_n)$ of degree d. If the initial condition is the same for every server then the new mean field (Section 2.3.2) is given by $\phi^{(n)}_{(k,s)} = \psi_s$ for any server $k \in V_n$ and queue length $s \in \mathcal{S}$ where ψ is the unique solution of the system of ODEs*

$$\frac{d\psi_s}{dt} = \mu(\psi_{s+1} - \psi_s \mathbb{1}_{\{s>0\}}) - \lambda \left(\psi_s(\tilde{g}_s + \tilde{g}_{s+1}) - \psi_{s-1}(\tilde{g}_{s-1} + \tilde{g}_s)\mathbb{1}_{\{s>0\}}\right), \tag{4.6}$$

with $\tilde{g}_s = \sum_{s' \geq s} \psi_{s'}$ and the shared initial condition.

The proof can be found in Appendix A.2.2. We give a brief sketch of the proof although the idea is similar to Theorem 3 and likewise follows proof plan A.2. We argue that we have exchangeability between all servers since they start from the same initial condition and their evolution is entirely dependent on their local neighbourhoods which all look the same since we have a regular graph. Using this exchangeability we again take a classical mean field approach to look at proportions of servers with a given queue length. Through this, the sums over the neighbourhoods cancel out leaving the form of the classical mean field.

We now give a few remarks regarding Theorem 5.

Remark 9 (Independence of n and d)**.** The system of ODEs (4.6) is independent of both the system size n and graph degree d. This means the solution will also be independent of n, d. Thus the mean fields of two arbitrary servers in different systems on regular graphs will both be the unique solution of (4.6) and so equal. However, we do not expect the true queue length distributions to be equal. A denser graph with higher degree d should be better at evenly distributing arrivals and so have a different queue length distribution. This means that on regular graphs, the new mean field approximation is not capturing the differences in graph structures.

Remark 10 (Fully connected graph)**.** We note that the fully connected graph is also a regular graph and so Remark 9 suggests that the mean field on any regular graph is the same as the fully connected graph. Like in the proof of Corollary 3, if we sum (4.6) over s we find the ODE of the classical mean field on the homogeneous fully connected graph (2.4). The new mean field on regular graphs coincides with the classical mean field on the fully connected graphs and so this suggests that there also exist scaling like interpretation of the new mean field. We explore this in the following section.

Remark 11. Finally, we explain why both JSQ(2) policies with and without replacement yield the same result. We notice in the proof of the proposition the neighbourhoods cancel out and so the same happens if we include a server in its own neighbourhood in the case of replacement.

4.4.2 Scaling results

In Section 4.4.1 we discussed the particle-like new mean field of finite systems. In Remark 10 it was noted that the new mean field result from Theorem 5 is the same on all regular graphs including the fully connected graph. Given the multiple mean field interpretations of the classical mean field on the homogeneous fully connected graph (Section 2.3.1) we wish to find similar scaling limit interpretations for the new mean field. We obtain a similar result to Theorem 4 and compare it to existing work [12] which finds fluid limits for growing sequences of systems on graphs.

We begin by constructing a sequence of systems. Let $(n(a))_{a\in\mathbb{N}}$ be an integer sequence $n(a) \to \infty$ as $a \to \infty$ and $(d(a))_{a\in\mathbb{N}}$ an increasing integer sequence where $d(a) = O(a^p)$ for some $p \in (0, 1]$ (the weaker condition of $d(a) \to \infty$ should be sufficient here Remark 8). Then we construct a sequence of systems $(X^{(n(a))})_{a\in\mathbb{N}}$ according to Section 4.1, where the a-th element has $n(a)$ servers and a regular graph structure $G_{n(a)}$ of degree $d(a)$.

Definition 6 (Occupancy state)**.** Consider a system $X^{(n)}$ according to Section 4.1 with regular graph structure $G_n = (V_n, E_n)$ of degree d. Then let $x_s^{(n)}(t)$ be the proportion of servers in the system that have queue length s at time t which is given by

$$x_s^{(n)} := \frac{1}{n} \sum_{k\in V_n} X^{(n)}_{(k,s)}(t).$$

Then we define the full occupancy state $x^{(n)}(t) := (x_s^{(n)}(t))_{s\in\mathcal{S}}$.

Notice that Definition 6 is almost the same as Definition 5 if we only consider 1 group. Similarly we note that $x_s^{(n)}(t) \in [0, 1]$ for all $n, d \in \mathbb{N}$ and $t \in \mathbb{R}^+$ and so we can compare this quantity for two systems of different sizes or graph structures. We look at the expectation of this quantity as $n \to \infty$ to obtain the LEOS like in Theorem 4.

Theorem 6 (LEOS)**.** *Consider the above sequence of systems $(X^{(n(a))})_{a\in\mathbb{N}}$ with regular graph structure $G_{n(a)}$ of degree $d(a) = O(a^p), p \in (0,1]$. Each system has occupancy state given by $x_s^{(n(a))}(t)$ for servers of queue length s (Definition 6). If each server has the same initial condition x, then the expectation of the occupancy state converges as such*

$$\mathbb{E}\left[x_s^{(n(a))}(t)\right] \to \psi_s(t,x), \quad \text{as } a \to \infty$$

where $\psi = (\psi_0, \psi_1, ...)$ is the unique solution of the system of ODEs

$$\frac{d\psi_s}{dt} = \mu(\psi_{s+1} - \psi_s \mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_s \tilde{g}_s - \psi_{s-1}\tilde{g}_{s-1}\mathbb{1}_{\{s>0\}}\right) \tag{4.7}$$

with $\tilde{g}_s = \sum_{s' \geq s} \psi_{s'}$ and initial condition x. We call $\psi(t,x)$ the LEOS of $(X^{(n(a))})_{a\in\mathbb{N}}$.

The proof of Theorem 6 follows the same procedure as Theorem 4 although a full proof is given in Appendix A.3.2 to address the details of how the sequence scales. Before we discuss the implications of Theorem 6 we give a similar but much stronger result [12, Theorem 2.1]. Given in our notation, they say that if the graph $G_{n(a)}$ of the a-th system in the aforementioned sequence $(X^{(n(a))})_{a\in\mathbb{N}}$ satisfies

$$1. \quad \min_{k\in V_{n(a)}} \left|\mathcal{N}_k^{(n(a))}\right| \to \infty \text{ as } a \to \infty \tag{4.8}$$

$$2. \quad \sum_{k_1 \in \mathcal{N}_k^{(n(a))}} \frac{1}{\left|\mathcal{N}_{k_1}^{(n(a))}\right|} \to 1 \text{ as } a \to \infty \text{ for all } k \in V_{n(a)} \tag{4.9}$$

then the occupancy state converges weakly to the deterministic fluid limit i.e. $x_s^{(n(a))} \to \psi_s$ as $a \to \infty$ where ψ is the unique solution of the system of ODEs (4.7).

We first discuss the differences in the conditions on graphs. Condition (4.8) on the growth of neighbourhoods is more relaxed than our $O(n^p), p \in (0,1]$, but as mentioned in Remark 8 only enforcing that all neighbourhood sizes tend to infinity as the system size grows to infinity should be sufficient. It is not trivial to interpret the class of graphs that satisfy Condition (4.9). We show that for finite n

$$\sum_{k_1 \in \mathcal{N}_k^{(n)}} \frac{1}{\left|\mathcal{N}_{k_1}^{(n)}\right|} = 1 \quad \text{for all } \ k \in V_n \quad \iff \quad G_n \text{ is regular.} \tag{4.10}$$

The forward direction is shown in the proof of Theorem 7 found in Appendix A.4 and the backwards direction can be seen since all neighbourhoods are of the same size. Therefore we suggest that Condition (4.9) equates to saying that the graph $G_{n(a)}$ is regular in limit. This is a weaker condition than ours which requires that all finite graphs $G_{n(a)}$ are also regular. We require this in order to apply Theorem 5 and use the 'independence of n' property in our proof (Appendix A.3.2). Although, this condition of regular graphs in limit is not particularly surprising given that the systems of ODEs (4.6) and (4.7) are the same and have unique solutions. Thus the fluid limit is the same as the mean field for a finite system on regular graphs (see Remark 12). In Section 4.4.3 we discuss whether the condition of regular graphs is necessary.

With regards to the limiting statement itself, [12, Theorem 2.1] is much stronger than Theorem 6. Due to our method using the new mean field error term we are restricted to saying that the expectation of the occupancy state converges to the unique solution of (4.7). Whereas [12, Theorem 2.1] makes a stronger statement by saying that the occupancy state itself converges weakly to the same quantity. They do this by using martingales to argue for existence of the limit and describing the fluid limit as an integral equation. It is not surprising that they converge to the same object. The fluid limit exists, thus we can take the limit inside of the expectation in Theorem 6 using the dominated convergence theorem and then since it is deterministic, the expectation should be the same as the fluid limit.

We now discuss the implications of Theorem 6 and [12, Theorem 2.1] to how we understand the new mean field, in the following remarks.

Remark 12 (Interpretation of mean field)**.** As an implication of [12, Theorem 2.1] we find that the fluid limit is the unique solution to the same system of ODEs (4.6) from Theorem 5. This means the fluid limit of a sequence of systems subject to Conditions (4.8), (4.9) is somehow 'equal' to the new mean field on a finite regular graph. It is difficult to see how to interpret this and we must be precise with the meaning of equal since the state spaces are different sizes. As per Theorem 5, the new mean field is the same for every server in the system and so we can just consider the mean field for one arbitrary server in a finite system on a regular graph structure given by $\phi^{(n)}_{(k,s)}$. Then $\phi^{(n)}_{(k,s)}$ is equal to the fluid limit of a growing sequence of systems where the graph is regular in limit. Therefore there exists a particle and scaling-like interpretation to the new mean field much like Remark 7.

Remark 13. We can find a corollary similar to Corollary 1 which says that for any finite system, we can construct a sequence of systems such that the fluid limit is equal to the particle like mean field of the finite system. This is not a particularly strong statement given the previous results of the section. Our aim is to find such a result for arbitrary graphs and service rates for which we do in Chapter 5.

4.4.3 Conditions on graphs

In Section 4.4.1 we saw that the mean field for an arbitrary server is insensitive to the size and degree of the regular graph. In Remark 10 we discussed that this notably means that the mean field of an arbitrary server on a regular graph is the same as the fully connected graph. An arising question is can we find all graphs such that the mean field is insensitive to the graph? What are the necessary conditions on the graph for the mean field to be the same on the fully connected graph? In this section we show that regular graphs are in fact also necessary for the mean field to be equal to a system on the fully connected graph.

Note: When we refer to 'the fully connected graph' we can choose a system of any size since it is independent of n, issues of this kind are discussed in Section 4.4.1. Likewise when we refer to two mean fields being equal we mean for any two arbitrary servers. These details are omitted when unimportant.

Theorem 7. *Consider a system $X^{(n)}$ as described in Section 4.1 with graph structure G_n and new mean field $\phi^{(n)}$ with initial condition x for each server. Further consider the mean field of a system on the fully connected graph given by ψ unique solution of (4.6) with the same initial condition x. Then*

$$\phi^{(n)}_{(k,s)} = \psi_s \quad \textit{for all} \quad k \in V_n \quad \iff \quad G_n \ \textit{a (disjoint union of) regular graph(s).}$$

The full proof can be found in Appendix A.4. The backwards direction follows directly from Theorem 5. In the forwards direction we say the drifts are equal and derive a condition on the reciprocals of neighbourhood sizes from $g_{(k,s)}$. Then we prove by induction that only disjoint unions of regular graphs satisfy this condition. We now discuss various remarks regarding Theorem 5.

Remark 14. As remarked in Section 4.4.2 we find Condition 4.10 in the proof of Theorem 5. The restriction on the sums of reciprocals of neighbourhood sizes is rather strong because the system of equations are very dependent. Not only do the same terms appear in the equations centred around all their neighbours, but also the number of terms in an equation is dictated by neighbourhood size.

Remark 15. We note that the appearance of disjoint unions of regular graphs is expected given that it holds for regular graphs. We see that a system on a disjoint union of p many regular graphs has an arrival rate of λn_i to the i-th disjoint subgraph where n_i are the number of servers in the subgraph. This is equivalent to having p many smaller systems each with regular graph structures of size n_i and each independent since they are disjoint. Then even if they have a different degree, their mean fields are still the same by Theorem 5.

Remark 16. Another way to view Theorem 7 is that regular graphs form a complete equivalence class with respect to the mean field. Theorem 7 does not require any sequences of systems however there are two ways of interpreting this equivalence class. The particle-like perspective is as follows: any two arbitrary servers from two systems in the class have the same mean field but a different error term to their queue length distribution based on their size n and degree d. The scaling limit approach would suggest that if we take some system from the class and construct a sequence whereby it is scaled up with $|\mathcal{N}_k^{(n)}| \to \infty, n \to \infty$ neighbourhoods, then the fluid limit (or LEOS) of this sequence is the same for any starting system (Remark 12, 13).

Remark 17. In this section and the rest of Section 4.4 we have discussed classes, sequences and limits of various systems and we wish to describe how they all fit together. For this we have a visualisation (Figure 4.3) of the sequence space of systems described in Section 4.1 with the same arrival rates and processing rates λ, μ respectively. We discuss the particle-like mean field for finite systems along the sequence and if applicable, the fluid limit (or LEOS) of the sequence. We describe each of the groups below:

1. The space of all sequences of homogeneous systems described in Section 4.1 each with arrival rate λ and processing rate μ.

2. This is the space of all sequences with graphs where the neighbourhood sizes tend to infinity and thus Proposition 1 holds. This suggests that as the systems grow large, the error term in Proposition 1 vanishes. Although, we cannot say whether the mean

field 'converges' or if the fluid limit exists. Note that $O(n^p)$ could be replaced with the weaker condition $|\mathcal{N}_k^{(n)}| \to \infty, n \to \infty$ (Remark 8).

3. This is the space of all sequences that satisfy Conditions 4.8, 4.9 from [12]. For these sequences we know that the fluid limit exists and is given by the same system of ODEs as the particle like new mean field of any finite system on a regular graph [12, Theorem 2.1]. Note that $O(n^p)$ could be replaced with the weaker condition $|\mathcal{N}_k^{(n)}| \to \infty, n \to \infty$ (Remark 8).

4. This is the space of all sequences such that every system in the sequence has a regular graph (not necessarily of the same degree). We therefore know that the mean field of an arbitrary server in each system is the same for every system in the sequence and also between sequences (Theorem 7).

5. This is the intersection of 2. (or 3.) and 4. and therefore in addition to their properties we also know that the LEOS exists and is the same as the fluid limit in 2. (Theorem 6).

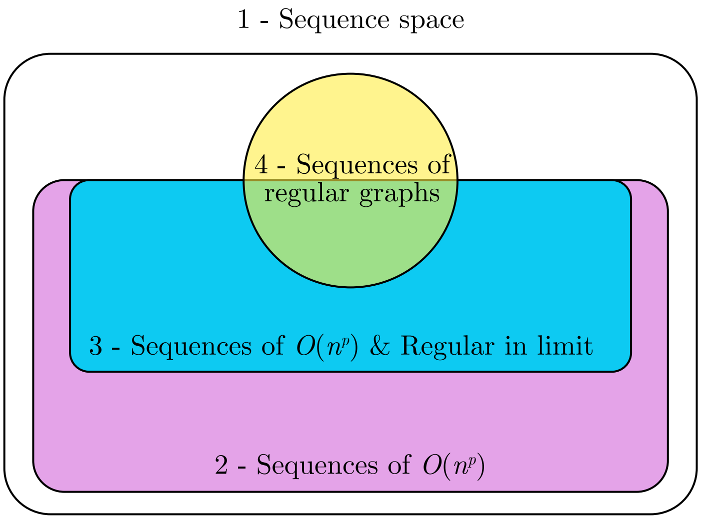

Figure 4.3: Visualisation of the sequence space of all systems.

4.5 Pair approximation

As mentioned in Section 2.1, Gast developed the pair approximation for homogeneous systems on regular graphs in [11]. The motivation for this was the observed poor performance of the mean field on these graphs. As discussed in Section 4.4.2, the mean field of homogeneous systems on regular graphs of any degree is the same as on the fully connected graph, thus the mean field does not capture the restrictions imposed by the graph and so performs poorly. We give a very brief heuristic comparison of the interpretation of the mean field and pair approximation as a way to convey the short comings of the mean field on graphs.

The pair approximation is constructed by considering the proportions of pairs of servers that have queue lengths (i, j). One then considers the evolution of a pair of servers with state (i, j). The state becomes $(i - 1, j)$ when there is a departure from the server of queue length i at rate μ and the state becomes $(i+1, j)$ when there is an arrival to server of queue length i. Just as with the mean field, the arrival rate is more difficult to compute. There can either be an arrival to this pair, or to any of the other pairs involving the server of queue length i. Thus the next stage of the analysis is to consider the proportions of all connected triples that have queue length (l, i, j). Gast then explains that the process in terms of triples is not density dependent (since the states are pairs) and so the idea is to replace the dependency on proportion of triples by proportions of pairs. This makes the arrival rate to the pair depend on the proportions of pairs (rather than triples) and in so doing, become a density dependent process which allows for the construction of an ODE to describe the evolution (see Section 2.3.1 for more explanation on the population process).

In a non-complete graph, although two servers may not have an edge connecting them, the effect of an arrival or departure from one has a knock on effect passing through the graph to the other server. The mean field only considers the immediate neighbours of a servers and therefore does not capture these 'long range' dependencies. This is fine in the complete graph, but in sparse graphs this becomes a more significant issue. The pair approximation improves on this by considering the triples which expands its direct dependencies thus capturing the structure of the graph and performing better than the mean field.

4.6 Numerical Results

In this section we briefly discuss numerical results to illustrate some of the theoretical results of the chapter through examples. General comments on numerics such as discussion of buffer size can be found in Section 3.4. Consider a homogeneous system with arrival rate $\lambda = 1$ and processing rate 1.5. We consider systems of $n = 10, 20, 30, 40$ servers on regular rings of degree $2, 4, 6, 8$ respectively. See Figure 4.2 for an example of the two-ring. In Figure 4.4 we provide the average queue length against time for the simulated system and computed mean field. We simulate each system 2500 times and the mean result is given with 95% Wald confidence interval.

We have very similar observations to the numerical results in Section 3.4. The mean field curve is the same for all systems which is expected by Theorem 5. Moreover the error between the mean field and simulation behaves in an $O(1/n)$ fashion as per Theorem 2 because we

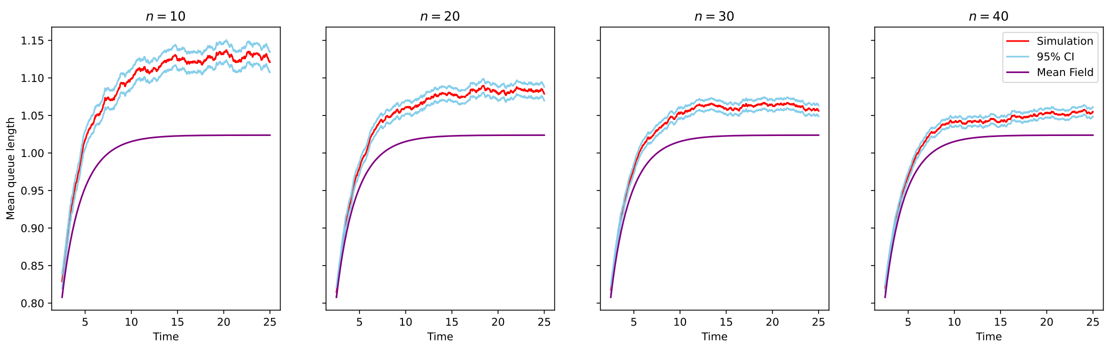

Figure 4.4: Average queue length of simulation vs mean field approximation for a proportionally scaled system.

have growing neighbourhoods of size $n/5$. Furthermore the $O(1/n)$ error appears to have a consistent constant across all the systems such that the decrease is uniform even with small n (see Section 3.4 for more details).

Chapter 5

Heterogeneous service rates and graph structures

The main focus of Chapter 3 was systems with fixed proportions of heterogeneous rates on the fully connected graph. Then in Chapter 4 we looked at homogeneous systems on regular graphs. We showed that for specific constructions the particle-like new mean field is independent of n. This was then used to find the limit of the expected occupancy state (LEOS). In this chapter we combine the concepts from Chapters 3 and 4 to study heterogeneous systems on graph structures which have received relatively limited attention in the literature [15]. We generalise the previous results on the independence of n and also answer the question posed at the end of Section 2.3.3 and find a type of growth such that particle-like mean field is is equal to the LEOS of a sequence with said growth.

5.1 Model, drift and mean field

We use the same model outlined in Section 2.2. This is an extension of the specific case described in Section 4.1, except we replace the homogeneous processing rate of server k with a heterogeneous rate μ_k. For completeness we provide a brief description. We have a system of n servers that exist on the vertices of a graph $G_n = (V_n, G_n)$. We then have Poisson arrivals at rate λn and server $k \in V_n$ has exponential processing times with heterogeneous processing rates μ_k. For assignment of arriving jobs we apply JSQ(2) on graphs without replacement. We note that like in Chapter 3 and unlike Chapter 4 we do not obtain the same results for both with and without replacement. We provide brief results on replacement but without replacement is the leading policy. We denote the model by the (time homogeneous) continuous time Markov process $X^{(n)}$ with notation given by Section 2.3.1.

The transition rates and derivation can be taken from Section 4.1, namely (4.3) and (4.4), with the small amendment of changing the homogeneous processing rate to the heterogeneous one described above. This yields the following transition rates

$$X^{(n)} \to X^{(n)} - e^{(n)}_{(k,s)} + e^{(n)}_{(k,s-1)} \quad \text{at rate } \mu_k X^{(n)}_{(k,s)} \mathbb{1}_{\{s>0\}} \tag{5.1}$$

$$X^{(n)} \to X^{(n)} + e^{(n)}_{(k,s+1)} - e^{(n)}_{(k,s)} \quad \text{at rate } \lambda X^{(n)}_{(k,s)} (g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}). \tag{5.2}$$

The definition of $g^{(n)}_{(k,s)}$ is unchanged from Section 4.1 since this only pertains to the graph structure, hence recall that

$$g^{(n)}_{(k,s)} := \frac{1}{2} \sum_{k_1 \in \mathcal{N}_k} \sum_{s_1=s}^{b} X_{(k_1,s_1)} \left(\frac{1}{\left|\mathcal{N}^{(n)}_k\right|} + \frac{1}{\left|\mathcal{N}^{(n)}_{k_1}\right|} \right).$$

Unsurprisingly we also find the same drift (4.5) from Section 4.2 with the same change to heterogeneous transition rates. For a state $x^{(n)} \in \mathrm{conv}(\mathcal{X}^{(n)})$ given by

$$f^{(n)}_{(k,s)}(x^{(n)}) = \mu_k x^{(n)}_{(k,s+1)} - \lambda x^{(n)}_{(k,s)}(g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}) - \left(\mu_k x^{(n)}_{(k,s)} - \lambda x^{(n)}_{(k,s-1)}(g^{(n)}_{(k,s-1)} + g^{(n)}_{(k,s)})\right) \mathbb{1}_{\{s>0\}}. \tag{5.3}$$

Finally we obtain the same conditions from Section 4.3 on sequences of systems for the growth of neighbourhoods if we wish to apply Proposition 1 and thus for a sequence of systems all neighbourhood sizes must grow with $O(n^p)$, $p \in (0, 1]$ (or the more relaxed condition of Remark 8).

5.2 Regular graphs

We now discuss the most direct extension of Theorems 3 and 5. Consider a system on a regular graph G_n of degree d, with m different types of servers $V^{(n)}_l$. Two servers, k, k' are both of type l if they both satisfy the following two conditions

1. the same processing rate $\mu_k = \mu_{k'} = \mu^l$ (5.4)
2. the same proportions of types of servers in their neighbourhood. (5.5)

This definition is recursive and thus difficult to immediately understand. We try to illustrate this by first considering the case where we assume all servers of the same processing rate are also of the same type. In this case Condition 5.5 means that every server of the same processing rate, must be connected to the same proportions of servers of other processing rates. We give examples of systems where a type has a unique processing rate in Figures 5.1, 5.2 and 5.3. If we do not assume that every server with the same processing rate is the same type, then we can obtain systems such as in Figure 5.4. We remark that the conditions on type are quite strong and encode a lot of information regarding how the processing rates are distributed across servers. It is also difficult to verify the type of a server without verifying the type of all servers in a graph.

Notice that the conditions on type of servers refer to proportions and thus two servers from different sized systems could be the same type. Let α^j_l be the proportion of type j server in the neighbourhood of type l servers. Define $\mathcal{L}_k$ to be the set of all types of servers in the neighbourhood of server k (which is the same for servers of the same type). We then obtain the following theorem.

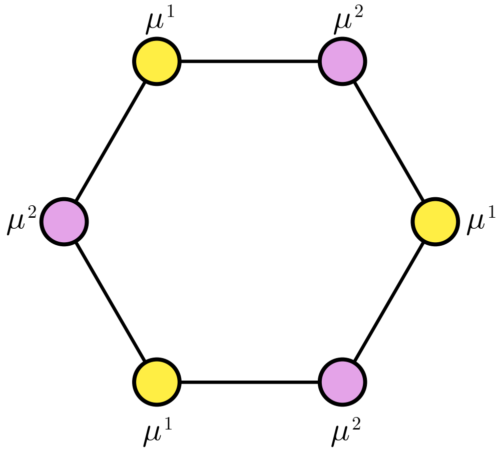

$d = 2,\ m = 2$

Figure 5.1: Regular graph of 6 servers of degree 2 with two types of servers denoted by colour and unique processing rates. All pink type servers are connected to only yellow type, all yellow type only connected to pink type.

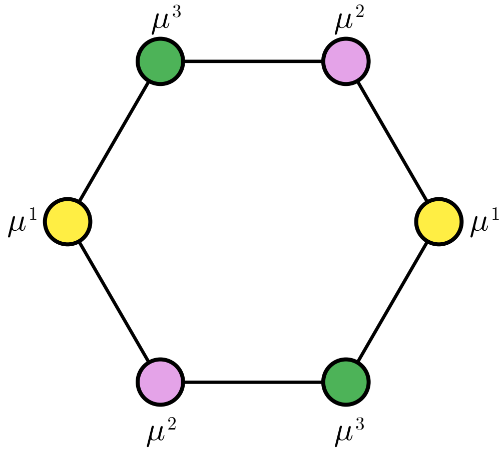

$d = 2,\ m = 3$

Figure 5.2: Regular graph of 6 servers of degree 2 with three types of servers denoted by colour and unique processing rates. Each colour is connected to both other colours thus forming three types of servers.

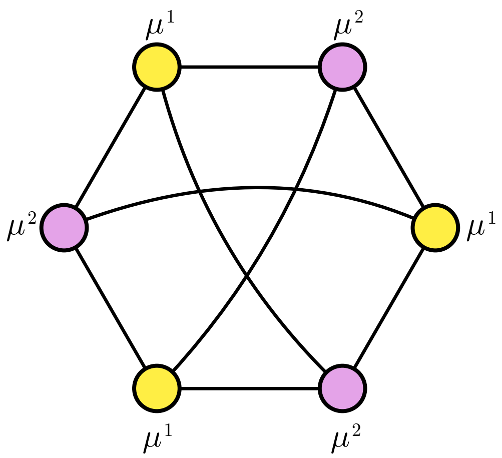

$d = 3,\ m = 2$

Figure 5.3: Regular graph of 6 servers of degree 3 with two types of servers denoted by colour and unique processing rates. All pink type servers are connected to only yellow type, all yellow type only connected to pink type.

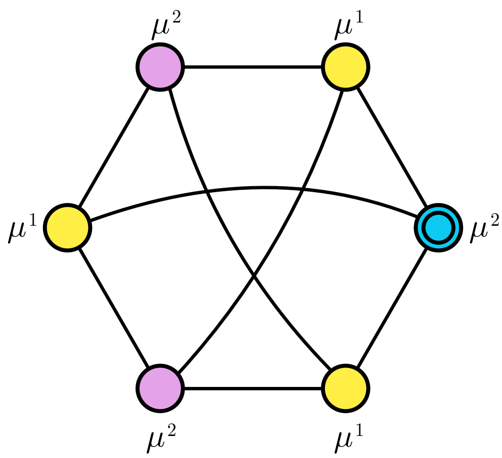

$d = 3,\ m = 3$

Figure 5.4: Regular graph of 6 servers of degree 3 with three types of servers denoted by colour but only 2 unique service rates. Blue is only connected to yellow, which are each connected to blue and two pink. Pink are only connected to yellow.

Theorem 8. *Consider system $X^{(N)}$ on regular graph G_N of degree d with m types of servers. Then the mean field for server $k_1 \in V_l^{(N)}$ (with initial condition x equal for all servers of the same type) is given by $\phi^{(N)}_{(k_1,\cdot)}(x,t) = \psi_{(l,\cdot)}(z,t)$ where ψ is the solution of the system of ODEs*

$$\frac{d\psi_{(l,s)}}{dt} = \mu^l(\psi_{(l,s+1)} - \psi_{(l,s)}\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_{(l,s)}(g_{(l,s)} + g_{(l,s+1)}) - \psi_{(l,s-1)}(g_{(l,s-1)} + g_{(l,s)})\mathbb{1}_{\{s>0\}}\right) \tag{5.6}$$

where $z = (x_1, ..., x_m)$ (x_l condition for type l servers) and $g_{(l,s)} = \sum_{j\in\mathcal{L}_l}\sum_{s_1\geq s}\alpha_l^j\psi_{(j,s)}$.

The proof is found in Appendix A.2.3 and is very similar to that of Theorems 3 and 5 and follows proof plan A.2. We now consider a few remarks regarding this theorem.

Remark 18 (Independence of n, d)**.** Much like for Theorems 3 and 5, the system of ODEs (5.6) is independent of n, d, up to the degree satisfying the conditions on proportion of type. It is not obvious how one could in general scale the system given the strong conditions on type, but for example if n grows by a factor a and the degree grows by a factor a then we could construct it accordingly. What the result means is that fixed proportions of types in the neighbourhoods preserve the heterogeneity with respect to the mean field. As discussed in Section 4.5, the mean field only directly depends on its neighbours and thus the type condition on the neighbourhood is sufficient.

Remark 19 (JSQ with and without replacement)**.** We find that Theorem 8 holds whether we use JSQ(2) with replacement or not since the definition of type only refers to proportions of the neighbourhood. In the case of replacement a server is in its own neighbourhood and so must be considered for the proportions. This does mean that we could have two systems using different policies have the same mean field.

Remark 20. If we consider the fully connected graph with replacement and m different groups of the same rate, then any two servers of the same rate are the same type. If we have $N_2 = aN_1$ then we obtain Theorem 3. Likewise if we consider a homogeneous regular graph then clearly all servers are of the same type and thus we obtain Theorem 5.

Remark 21. It should be possible to extend this result to nonregular graphs. We have not studied this since the definition of type becomes more cumbersome as it also looks at neighbourhood size. Thus the processing rates, proportions of other types and the sum of the reciprocal neighbours' neighbourhood sizes (proof A.4 of Theorem 7) must be equal. The conditions are strict and recursive and it becomes difficult to imagine nonregular solutions although they likely exist.

5.3 Arbitrary graphs

In this section we consider heterogeneous systems on graph structures without restrictions on either. We introduce a type of scaling we call cluster scaling to obtain an 'independence of n' and a LEOS results. The idea is to take any system and replace the servers (nodes) of the graph by clusters of servers with the same processing rate. We define a Cluster Graph in Definition 7 and give two examples in Figures 5.5 and 5.6. This concept is similar to that of 'interference graphs' in [21, Section 3.1]

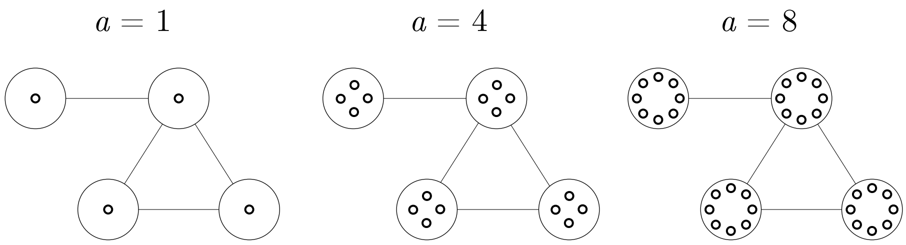

Figure 5.5: On the right and middle we have the degree 4 and 8 Cluster Graphs of the seed graph on the left. Interpret an edge between clusters as there existing an edge between all servers in each cluster.

Definition 7. Consider a graph $G_N = (V_N, E_N)$ where we replace the vertices with clusters of a many nodes. We then construct a graph G^c_{aN} by saying there exists an edge between any two nodes if there is an edge between their respective clusters. We call the graph G^c_{aN} a Cluster Graph of degree a generated by G_N.

Remark 22. We note that under Definition 7, nodes in the same cluster are not connected, which represents JSQ(2) without replacement. In the case of JSQ(2) with replacement, nodes in the same cluster should be connected. Unless stated otherwise we assume JSQ(2) without replacement and thus nodes in the same cluster are not connected. However, in the case of a loop in the graph and a server is connected to itself, all the servers in the corresponding cluster are also connected.

Using Definition 7 we can define a cluster scaled system.

Definition 8 (Cluster scaling)**.** Let $X^{(N)}$ be a system of N servers, with heterogeneous processing rates μ_k and graph structure G_N. Then the degree a Cluster Scaled System $\tilde{X}^{(aN)}$ of $X^{(N)}$ has Cluster Graph structure G^c_{aN} generated by G_N (Definition 7) where any server $k_1 \in V^{(aN)}_k$ (the k-th cluster) has processing rate $\mu_{k_1} = \mu_k$.

Similarly to Sections 3.3 and 4.4 we first look at finite cluster scaled systems and then the LEOS of a sequence of scaled systems.

5.3.1 Finite systems

For now we consider cluster scaling one arbitrary finite system and study the particle-like mean field. Similarly to Sections 3.3.1, 4.4.1 and 5.2 we look to show that the mean field for like servers in different sized Cluster Graphs is the same, thus obtaining a kind of 'independence of n' result. This yields the following theorem

Theorem 9. *Consider a heterogeneous system $X^{(N)}$ of N servers with graph structure G_N and its cluster scaled system $\tilde{X}^{(aN)}$ of degree a for some $a \in \mathbb{N}$ (Definition 8). Let $X^{(N)}$ have initial condition x and $\tilde{X}^{(aN)}$ initial condition z where for $k_1 \in \tilde{V}^{(aN)}_k$ (the k-th cluster), $z_{k_1} = x_k$. Then the mean field for any server $k_1 \in \tilde{V}^{(aN)}_k$ is equal to the mean field of server $k \in V_N$ i.e.*

$$\tilde{\phi}^{(aN)}_{(k_1,s)}(z,t) = \phi^{(N)}_{(k,s)}(x,t) \quad \text{for all } s \in \mathcal{S},\ t \in \mathbb{R}^+$$

where $\phi^{(N)}$ is the new mean field of $X^{(N)}$ outlined in Section 2.3.2.

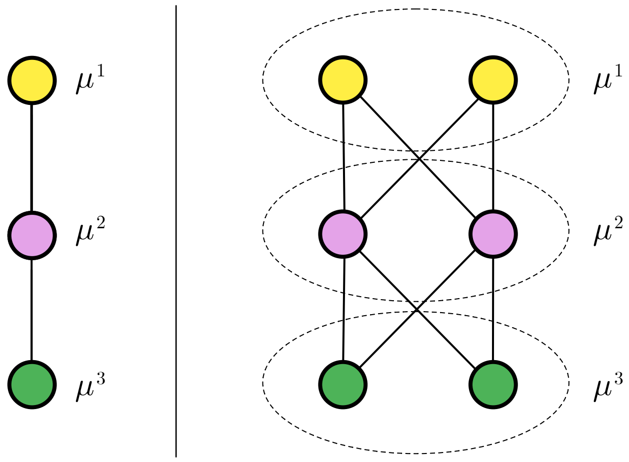

Figure 5.6: On the right we have a degree 2 Cluster Graph of the graph on the left. All inter-cluster connections are explicitly drawn. Note that the heterogeneity is preserved within a cluster.

The proof is similar to that of previous results Theorems 3, 5 and 8 and follows the proof plan A.2. However the full proof can be found in Appendix A.2.4. We now briefly discuss some remarks.

Remark 23 (Independence of n)**.** We yet again obtain an 'independence of n' result suggesting that for any two servers in different systems but from the same cluster have the same mean field. Just as seen in Theorem 3, this expresses that the heterogeneity and structure of the system is preserved with cluster scaling. However, as noted in Remark 18, the mean field only needs the structure to be preserved locally.

Remark 24. Theorem 9 is different to Theorem 8. If our initial system is a regular graph then all Cluster Graphs will also be regular, however the scaling is very specific and so we cannot obtain any regular graph from another. For example Theorem 8 allows for two regular graphs of the same degree but different sizes which Theorem 9 does not allow for since the degree of the cluster scaled graph with be increased by a factor a.

Remark 25. Theorem 9 holds for either JSQ(2) policy (with or without replacement) as long as the Cluster Graph (Definition 7) is chosen accordingly (see Remark 22). If we initially choose the fully connected graph and then JSQ(2) with replacement in Definition 8 then we obtain proportional scaling from Section 3.3.2 (Definition 4). Therefore under this scheme we also obtain Theorem 3 from Theorem 9.

5.3.2 Scaling results

In Section 5.3.1 we discussed scaling finite systems and their particle-like mean fields. Our goal is to now investigate sequences of cluster scaled systems for any starting system. We study the occupancy state for each cluster and use this to find the 'LEOS'. The purpose is to show that for an arbitrary heterogeneous system on graphs, we can interpret the mean field in either a particle or scaling sense. nonregular We follow a similar process to Sections 3.3.2 and 4.4.2. Consider a system $X^{(N)}$ of N servers with graph structure $G_N = (V_N, E_N)$

and heterogeneous rates μ_k for $k \in V_N$. Let $(d(a))_{a\in\mathbb{N}}$ be a sequence of integers such that $d(a) \to \infty$ as $a \to \infty$ with $d(1) = 1$ (the weaker condition Remark 8). We construct a sequence $(X^{(d(a)N)})_{a\in\mathbb{N}}$ where the a-th element is the degree $d(a)$ cluster scaled system of initial system $X^{(N)}$ (Definition 8). Note that we therefore have N many clusters in each system. We then construct the occupancy state of a cluster $k \in V_N$ in the $a-$th system in the following definition.

Definition 9 (Occupancy state)**.** Consider the a-th system from the above described sequence, $X^{(d(a)N)}$, then the set of servers in the k-th cluster is given by $V_k^{(d(a)N)}$ of size $d(a)$ in which all servers share processing rate μ_k (Definition 8). Then let $x_{(k,s)}^{(d(a)N)}(t)$ be the proportion of servers in the k-th cluster that have queue length s at time t which is defined as

$$x_{(k,s)}^{(d(a)N)}(t) := \frac{1}{d(a)} \sum_{k_1 \in V_k^{(d(a)N)}} X_{(k_1,s)}^{(d(a)N)}(t).$$

Then we define the full occupancy state as such $x^{(d(a)N)}(t) := (x_{(k,s)}^{(d(a)N)}(t))_{k=1,..,N,\ s\in\mathcal{S}}$.

This definition is similar to Definition 5 and 6 but importantly different because every cluster is the same size. Since the size of the cluster is $d(a)$ we see that $x_{(k,s)}^{(d(a)N)}(t) \in [0,1]$ and thus we can attempt to study the quantity as $a \to \infty$. For a slightly more detailed discussion of the occupancy state see Section 3.3.2. We now use the occupancy state for each cluster to find a LEOS for said cluster in the following theorem.

Theorem 10 (LEOS)**.** *Consider a sequence of cluster scaled systems $(X^{(d(a)N)})_{a\in\mathbb{N}}$, where each system has occupancy state given by $x_{(k,s)}^{(d(a)N)}$ for servers in cluster k with queue length s (Definition 9). If each server in the same cluster for any system in the sequence have the same initial condition x_k, then the expectation of the occupancy state converges as such*

$$\mathbb{E}\left[x_{(k,s)}^{(d(a)N)}(t)\right] \to \phi_{(k,s)}^{(N)}(t,x), \quad \text{as } a \to \infty$$

where $\phi^{(N)}$ is the new mean field of the initial system $X^{(N)}$. We can also refer to $\phi_{(k,\cdot)}^{(N)}$ as the LEOS of the sequence $(X^{(d(a)N)})_{a\in\mathbb{N}}$.

The proof of Theorem 10 follows the same process as Theorems 4 and 6, although a full proof can be found in Appendix A.3.3. In discussion of Theorems 4 and 6 in Sections 3.3.2 and 4.4.2 we have heavily discussed the meaning and strength of the LEOS. The aforementioned analysis is applicable here and as such we discuss only a few more remarks.

Remark 26. Notice that the neighbourhood size of a server in the k-th cluster of the a-th system is size $d(a)|\mathcal{N}_k^{(N)}|$. Therefore the growth rate of the neighbourhood sizes is exactly $d(a)$. Therefore $d(a) \to \infty$ as $a \to \infty$ is sufficient to apply Proposition 1 to the sequence of systems with the weaker condition of Remark 8.

The way to understand Theorem 10 is that the LEOS of a cluster scaled system is the same as the particle-like mean field of the initial system (seed system with $a = 1$). This is important since it means we can both view the mean field of the initial system as either a finite particle-like approximation, or a scaling-like limit of the expected behaviour. In Section 2.3.3 we discussed the differences between these interpretations but also observed that a scaling interpretation of the new mean field was not obvious. We posed the question of whether for a given finite system, there existed a type of growth such that the particle mean field is a scaling type mean field. To fully answer this question we rephrase Theorem 10 in the following corollary.

Corollary 4. *For any heterogeneous system $X^{(N)}$ of N servers on a graph structure G_N the new mean field is given by the LEOS of the sequence of cluster scaled systems $(X^{(d(a)N)})_{a\in\mathbb{N}}$ with $(d(a))_{a\in\mathbb{N}}$ any integer sequence such that $d(1) = 1, d(a) \to \infty$ as $a \to \infty$ (Definition 8).*

Proof. This follows directly from Theorem 10. □

We rephrase Theorem 10 into Corollary 4 to make it explicit that if we are given any fixed finite system, we can find a sequence such that the LEOS is the mean field. Thus finding a scaling-like interpretation to the new mean field.

We notice that in both Theorem 10 and Corollary 4 the specific sequence of cluster sizes is arbitrary as long as it starts at 1 and tends to infinity. There are infinitely many sequences that satisfy this condition. Therefore the way we interpret the cluster scaling is not as a specific scaling that depends on the specific cluster sizes but rather as a type of growth. The growth that is associated with any cluster scaling yields the LEOS and as such it preserves the structure of the initial system with respect to the mean field. When scaling homogeneous systems or proportionally scaled systems this is a property we take for granted. We explore how to interpret this as growth in Section 5.4.

In Section 3.3.2 we compared the strength of the statement made by a fluid limit and the LEOS, especially in comparison with [12]. We emphasise that the LEOS is not as strong since it is a limit of the expectation of the occupancy state rather than the weak limit of the occupancy process. Future work could aim to find the fluid limit of cluster scaled systems which would be necessary to make a rigorous connection between the new mean field and the scaling limit approach. Although this is beyond the scope of this Bachelor project. Perhaps our main contribution is how to frame the scaling limit problem for arbitrary systems in terms of cluster scaling and for this to be in such a way that is meaningful to find a fluid limit.

Remark 27. Akin to Remark 6, we notice that if the fluid limit exists then we can take the limit inside the expectation of the occupancy process by the dominated convergence theorem yielding the expectation of the fluid limit. Then since the fluid limit is deterministic we should have that the expectation is equal to the fluid limit. Hence with this rough reasoning we suggest if the fluid limit exists, it should be equal to the LEOS. We also saw that they are equal for homogeneous systems in Section 4.4.2 when comparing our LEOS with the fluid limit from [12].

5.4 Convergence issues

Throughout our work we have followed a pattern of constructing sequences of systems with very specific growth and then trying to express the behaviour of these systems as the number of servers grows large. Our main tool has been to find the limit of the expected occupancy state (LEOS). The occupancy state is the proportion of like servers that have a certain queue length. Defining this in general is not always easy for heterogeneous systems, like servers must have the same processing rates and so as the systems scale they must not introduce 'more heterogeneity'. This is a very limiting condition on the type of sequences we can consider.

The idea of the fluid limit is that the occupancy process converges weakly to some deterministic quantity and thus as $n \to \infty$ the process behaves deterministically. We now discuss some examples of sequences of systems that behave deterministically in limit and how this relates to the particle and scaling interpretations of the mean field and the proposed role of cluster growth. These arguments are not formalised nor rigorous and formalising them is beyond the scope of this Bachelor project. We want to focus on the concepts of how sequences tend towards deterministic behaviour.

In Theorem 10 we only found the LEOS for cluster scaled sequences but as discussed in Remark 27 we believe this should be equal to the fluid limit (if it exists). We continue by making the assumption that the fluid limit is equal to the LEOS. In Figure 5.7 we try to represent how a sequence of cluster scaled systems tends to the deterministic behaviour described by its fluid limit. On the y-axis we have the 'error w.r.t. limiting behaviour' - this is an abstract distance between the behaviour of a finite system in the sequence and the limiting behaviour of the sequence. In this case, it can be thought of as the error between the queue length distributions of servers in the same cluster in the finite and limiting system.

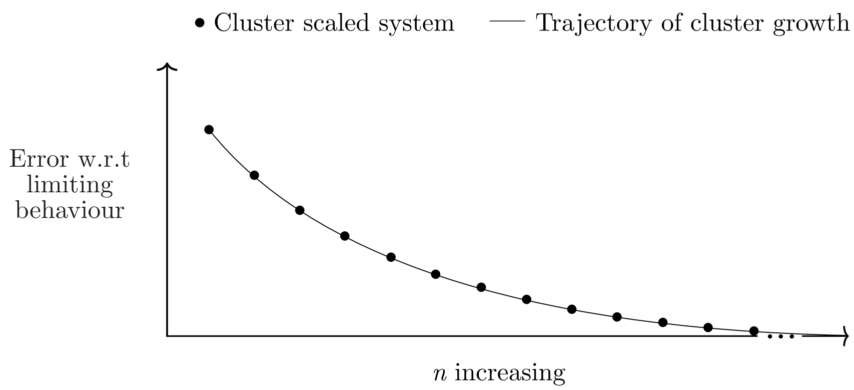

Figure 5.7: A sequence of cluster scaled systems with increasing degree against the error between the queue length distribution and the fluid limit. The trajectory line represents that any other cluster scaled sequence starting from the same system would fall on this line. The break in the x-axis represents n growing to infinity.

The points in Figure 5.7 show a specific cluster scaled sequence but as remarked in Section 5.3.2, there are infinitely many cluster scaled sequences with the same LEOS. Therefore we represent cluster scaling by a line which shows the trajectory of cluster growth. Furthermore notice that if we had any starting system along the curve, then the cluster growth would follow the same curve. This shows the self similarity between the structures of the cluster scaled systems.

In this example the limiting behaviour is the fluid limit of the cluster scaled sequence. We expect this to be equal to the LEOS and thus the particle-like mean field of the first system in the sequence (Theorem 10). From the opposite perspective, we know the mean field of a finite system is the LEOS of an implicit cluster scaled sequence. If we were to plot this implicit

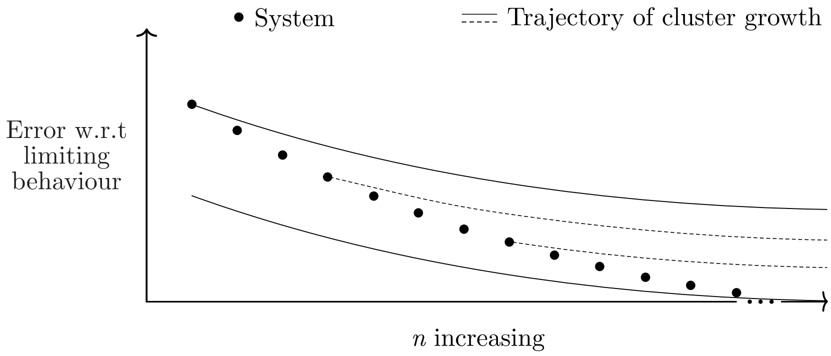

Figure 5.8: A sequence of systems with fast and slow servers but a fixed finite number more of slow servers than fast. The trajectories represent the path any cluster scaled sequence would take from any point on the curve. The break in the x-axis represents n growing to infinity.

sequence onto Figure 5.7 it would follow the trajectory of the black line i.e. the trajectory of cluster growth. The idea is that the mean field of a finite system is the 'limit' of the trajectory of cluster growth starting from that system. We explore this further in the next example.

We now look at a non-cluster scaled sequence. Consider a sequence of systems X that only consist of fast and slow servers but always contain exactly a fixed finite number of slow servers more than fast servers. We know that as n grows large, the impact of the finite extra slow servers disappears and the system behaves as a system with half slow and half fast servers. In Figure 5.8 the plotted points represent the aforementioned sequence X. As n grows large, the behaviour of the system tends to the deterministic fluid limit in which half of the servers are slow and half fast. This is the same fluid limit as a sequence of cluster scaled systems starting from a system with half slow and half fast servers (represented by the bottom trajectory in Figure 5.8).

Now suppose we find the (particle-like) mean field of a finite system $X^{(N)}$ along the plotted sequence X. By Theorem 10, the mean field is the LEOS of a cluster scaled system starting from the finite system $X^{(N)}$. The cluster scaling trajectory is represented by a dotted line in Figure 5.8. The mean field of $X^{(N)}$ is the 'limit' of the cluster scaling trajectory starting from the finite system $X^{(N)}$. By 'limit' it is meant that the mean field describes the limiting behaviour of a sequence along this trajectory. The error of this trajectory does not tend to zero because cluster scaling will maintain the imbalance of the finite extra slow servers and so this sequence's limiting behaviour is not the same as the sequence X. More simply, the mean field of any finite imbalanced system is not the same as the half and half system.

Consider a sequence of mean fields of all finite systems in the sequence $X^{(N)}$. This sequence should 'converge' to the mean field (fluid limit) of the half and half system. In this context converge should be understood as the mean fields for each group of like servers converges to avoid the issue of a growing state space. The observation from the last paragraph is that the mean field of a finite system is loosely the 'limit' of the cluster scaling trajectory starting from said finite system. Thus we can also loosely see our sequence of mean fields as the sequence of limits of cluster trajectories. Graphically this means the cluster growth trajectories in Figure 5.8 converge in error to the fluid limit trajectory.

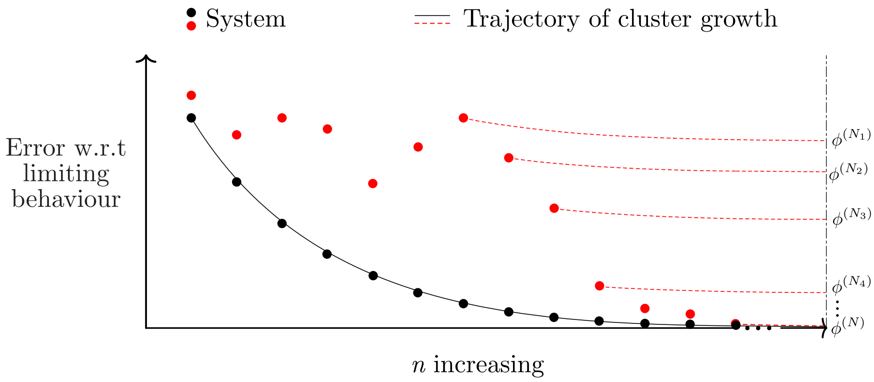

Figure 5.9: Sequence of homogeneous systems on the complete graph in black and sequence of systems with growing neighbourhoods that are regular in limit in red. The trajectories show cluster scaling starting from any point along the curve. The mean fields $\phi^{(N_i)}$ are the projected limits of the red dotted curves 'at infinity'. The error can be interpreted as the error between the occupancy state and fluid limit.

Lastly, consider a new sequence of homogeneous systems X on a sequence of graph structures that satisfy Conditions 4.8 and 4.9 meaning they have growing neighbourhoods and are regular in limit. These are the sequences studied in [12] and we discussed in Section 4.4.2. An example is plotted in red in Figure 5.9. The fluid limit of these sequences is the same as on the complete graph ([12, Theorem 2.1]) hence let Y be a sequence with the same arrival and processing rates as X but on the complete graph. The sequence Y is plotted in black dots in Figure 5.9. In this example we can interpret the 'error w.r.t. liming behaviour' as the error between the occupancy state of a system and the shared fluid limit of the sequences X, Y.

We can find the mean field for any finite system $X^{(N)}$ in the sequence X. As discussed above, we can loosely interpret this as the 'limit' of the cluster growth trajectory starting from $X^{(N}$ (represented by dotted red line in Figure 5.9). If we construct a sequence of the mean fields for each system in X then we see they 'converge' in error towards the fluid limit, which is the mean field of the fully connected graph. Again as mentioned above, this can be seen as the limit of the cluster scaling trajectories converging towards the fluid limit. In Figure 5.9 we see the the limits of the trajectories 'projected' at infinity, given by the mean fields. We then see these mean fields converge in error (downwards in the y-axis) towards the fluid limit of both X and Y.

5.5 Numerical results

We again provide some numerical results to elaborate on the examples in this chapter and visualise the theoretical results. We look at similar results to those discussed in Sections 3.4 and 4.6 and thus general comments from these sections also hold and are omitted or shortened here for brevity.

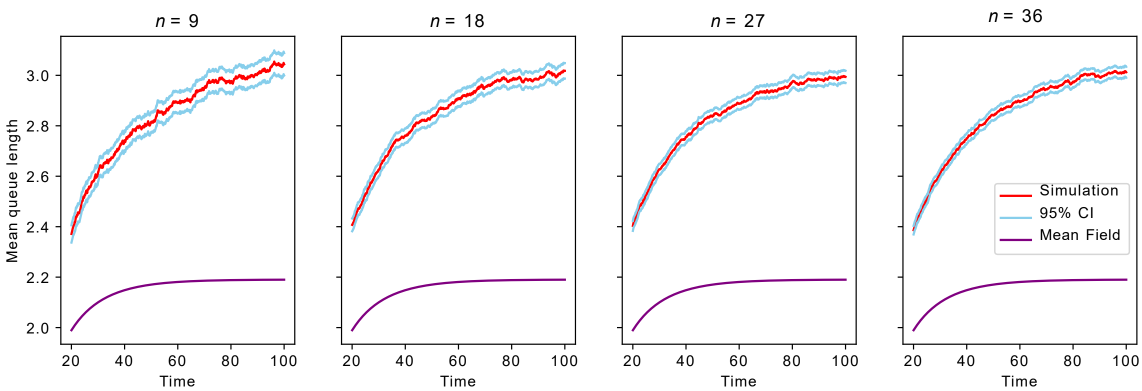

Figure 5.10: Average queue length of simulation vs mean field approximation for two-rings with three types of servers.

5.5.1 Regular graph

In our first example we consider a regular graph with heterogeneous rates. Namely we consider a growing two-ring with three different rates (and hence types) of servers. The visualisation in Figure 5.2 shows the system in the case $n = 6$, which grows with n by adding servers in the same pattern. We consider systems of size $n = 9, 18, 27, 36$ with arrival rate $\lambda = 1$, heterogeneous processing rates $\mu = 2, 3/2, 1/2$ and JSQ(2) without replacement. We have a buffer size of 20 such that we have negligible loss rate (see Section 3.4 for further discussion on buffer size).

Once again we look at the average queue length in Figure 5.10. The key difference between these systems and other numerical results is that we have not scaled the neighbourhood sizes. In Figure 4.4 we considered the homogeneous ring with growing neighbourhood sizes whereas in this example the neighbourhoods sizes are constant. This means that Proposition 1 does not hold which is clearly verified by the constant error between the simulation and mean field curves. We know from Theorem 8 that the mean field curve should be constant across the system sizes which is what we observe. The constant simulation line between the growing systems shows that they does not improve as it scales unlike our previous examples. A heuristic explanation is that the arrival rate to servers of the same type does not change because locally the neighbourhoods of a given type server look the same for any sized system and the arrival rate to the system scales with the number of servers in the system. Therefore, the probability to choose the neighbourhood of a given server is constant and then locally they look the same so the 'load balancing' is the same yielding a constant arrival rate. The point that we reiterate is that the vanishing error term in Proposition 1 is not due to the mean field approximation, but rather the systems improving with n until the servers behave independently as the neighbourhood sizes tend to infinity.

5.5.2 Cluster Graph

In this example we demonstrate the cluster scaling of the example given in Figure 5.5. We have arrival rate $\lambda = 1$, JSQ(2) without replacement and each cluster has processing rates $1, 2, 1.2, 1.1$ working clockwise from the top left. We consider the three systems with cluster sizes $a = 1, 4, 8$ and thus systems of size $4, 16, 32$. Figure 5.11 gives the average queue length

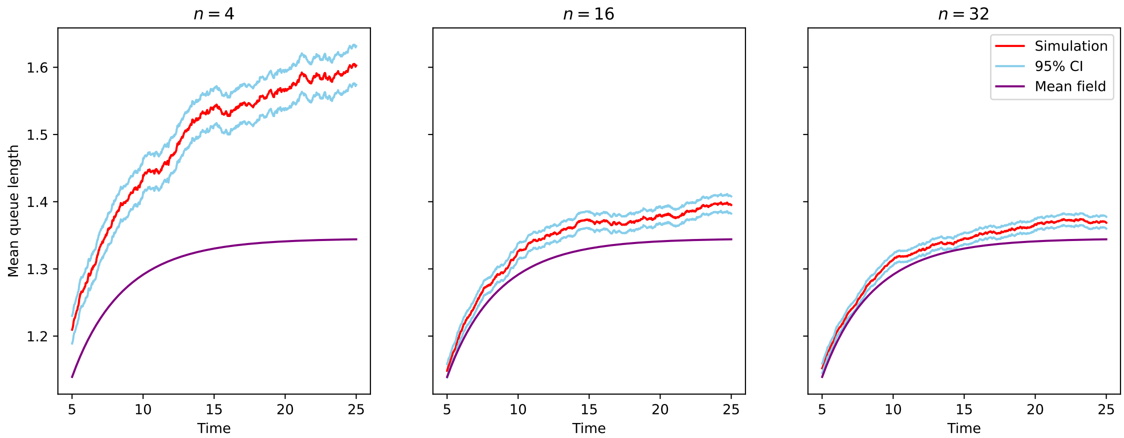

Figure 5.11: Average queue length of simulation vs mean field approximation for a cluster graph of four clusters of four rates for three degrees of scaling $a = 1, 4, 8$.

for the simulation and the mean field. The mean field curve is the same across the scaled systems as per Theorem 9. We also see the simulation curve tends to the mean field curve and the confidence intervals shrink as n increases. The error between simulation and mean field also appears to behave in an $O(1/n)$ sense, with no fluctuations due to the specific constants as described Section 3.4. This is likely due to the fact that by construction the cluster scaled systems are similar to the original system and thus the associated constants of the error term are similar. Similar observations were explained in more detail in Section 3.4.

5.5.3 Fast and slow servers

In the following example we expand on the convergence example from Section 5.4 with fast and slow servers. We have arrival rate $\lambda = 1$ and JSQ(2) with replacement on the fully connected graph. We have fast servers with processing rate 2 and slow servers with processing rate 1.2. In the system we always have four more slow servers than fast servers. The figure of mean queue length against time can be seen in Figure 5.12. Recall from the example in Section 5.4 that we expect that as n grows large, the impact of the extra four slow servers should go to zero and the system should behave as a proportionally scaled system from Chapter 3 with equal proportions of fast and slow servers. We see that the mean field curve appears to be converging to some curve as the proportional impact of the extra slow servers decreases and it behaves more like the system of half slow servers and half fast servers. This is better seen in Figure 5.12. Despite the mean field curve decreasing, we still see that the error between the simulation and mean field curves is decreasing in an $O(1/n)$ fashion which is exactly expected from Theorem 2.

In Figure 5.13 we show in orange the average queue length of the 'balanced' system with half servers of rate 1.2 and half servers of rate 2, whilst purple is the same mean field seen in Figure 5.12. This demonstrates the suggested convergence in behaviour between the system with four more slow servers than fast and the half and half system.

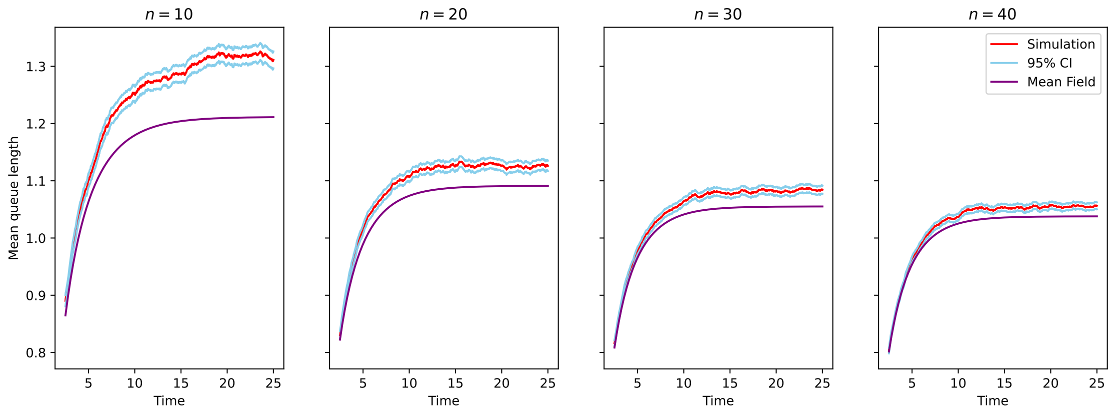

Figure 5.12: Average queue length of simulation vs mean field approximation for system with only fast and slow servers with exactly four more slow servers than fast.

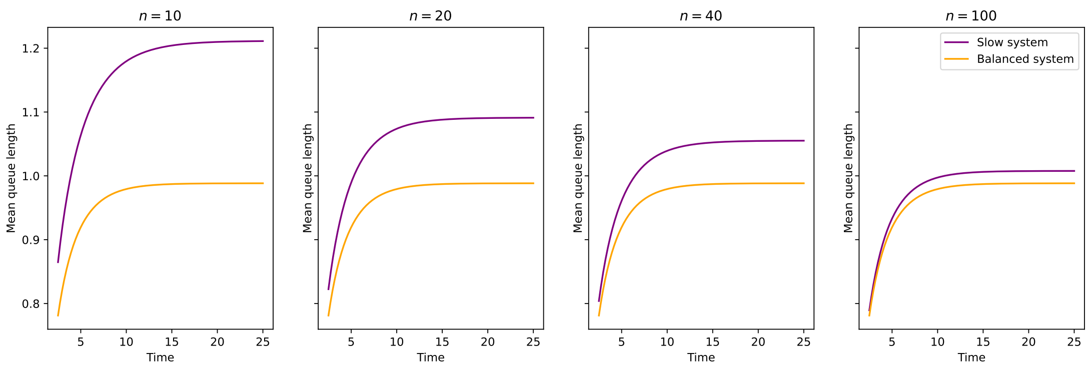

Figure 5.13: Average queue length of mean field for a balanced system with half slow and fast servers and the slow system with exactly four more slow servers than fast.

Chapter 6

Conclusions

In this work we studied the new mean field for multiple variations of the heterogeneous supermarket model on graphs. For each of the variations we obtained two main types of results. The first type considered scaling of finite systems such that the mean field is independent of the system size. The second type of results use the independence property to obtain a limiting result about the behaviour of a sequence of these scaled systems.

Independence of system size:
In Chapter 3 we considered systems on the complete graph with fixed proportions of service rates. We showed that the mean field of a proportionally scaled system is independent of the system size. This suggested that proportional scaling preserves the heterogeneity of the system. We note that we could take any system on the complete graph and proportionally scale it and so this gave initial insight to how we could scale general heterogeneous systems.

On the other hand, in Chapter 4 we looked at homogeneous systems on regular graphs. We showed that the mean field on any regular graph was independent of the size and degree. This importantly meant it was the same as on the fully connected graph and thus the mean field does not adequately capture the structure of the graph. We further discussed this in the context of the pair approximation. We also showed that for the mean field on some graph is equal to the mean field on the complete graph if and only if the graph is regular.

We combined the above results in Chapter 5 to obtain a result for heterogeneous systems with regular graphs. The result said that if the proportions of 'types' of servers are fixed then the mean field is independent of the size of the system or degree of the graph. The difficulty came in the recursive definition of the type of server which said that servers are of the same type if they have the same processing rate and connections to servers of other types.

These results showed firstly that mean field does not adequately capture the structure of graphs. Secondly that heterogeneous systems with the same proportions of rates preserve the heterogeneity. Finally it showed that the new mean field does not inherently depend on the system size as the notation suggests and that any observed dependence on system size is due to the structure or nature of the heterogeneity changing between different size systems.

Limiting results:
In Section 2.3.3 we discussed that the new mean field is interpreted as an approximation for

a finite system with some error term determining its accuracy. We noted that the classical mean field for homogeneous systems has an interpretation as a scaling limit as well as an approximation of a finite system. The question posed at the end of Section 2.3.3 was whether we could bridge this gap in interpretation for the new mean field.

In Section 5.3 we considered heterogeneous models with arbitrary graphs. Inspired by the way that proportional scaling in Chapter 3 preserved the heterogeneity of the system, we developed cluster scaling (of which proportional scaling is a special case of on the fully connected graph). In this scaling regime we take any system and replace its servers by clusters of servers with the same processing rate. We then say that two servers are connected if their clusters are connected. We then showed that the mean field of a server in a cluster was independent of the degree of scaling similar to the aforementioned 'independence of system size' results. Much like proportional scaling, cluster scaling preserves the heterogeneity even when increasing the system size.

In the next step we considered the proportion of servers in a cluster that have a given queue length called the occupancy state. Using the independence of the degree of scaling and the vanishing error term of Theorem 2 (or Proposition 1) we showed that the limit of the expected occupancy state (LEOS) for a cluster scaled system is exactly equal to the mean field of the pre-scaled system. The implication of this is that for any finite heterogeneous system with an arbitrary graph structure, we can either interpret the new mean field as an approximation of the system with some error term or as the limit of the expected occupancy state.

Further Study:
We discuss two avenues for further study. First of all, the LEOS is a limit of the expected occupancy state and therefore a much weaker limiting statement than the fluid limit which is the limit in distribution of the occupancy state. To fully show the scaling limit interpretation of the new mean field, the fluid limit for cluster scaling must be found. As noted in Remark 27, if the fluid limit exists then we expect this to be equal to the LEOS.

The other avenue is an extreme case of heterogeneity where the service rates are random. In this scheme the model itself is a random quantity and thus we have a kind of 'double randomness' whereby the evolution is random and the service rates are random. The model has probability distribution given by the joint distribution of the random service rates and the state space is all systems with rates in the support of the distribution of the random rates. The queue length probabilities and thus the mean field are also random variables with the aforementioned distribution. This immediately opens up questions such as: How can the mean field be interpreted? Should we consider the expectation of the mean field or queue length probabilities with respect to the joint distribution of the service rates?

Consider an example on the fully connected graph where the service rates are either slow or fast with probability 1/2. We expect that as n grows large there is some kind of 'law of large numbers' such that this system behaves the same as a proportionally scaled system with half slow and half fast servers. As discussed in Section 5.4, we expect this proportionally scaled system to behave deterministically as described by its mean field as n grows large and so does the random system also behave in the same manner? When is the mean field of a random model deterministic?

Bibliography

[1] 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. 1, 3, 4, 7, 8, 13, 19

[2] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich, "Queueing system with selection of the shortest one from two queues - an asymptotic approach," *Problemy Peredachi Informatsii*, vol. 32, no. 1, pp. 20–34, 1996. 1, 3, 4, 13

[3] S. Allmeier and N. Gast, "Mean Field and Refined Mean Field Approximations for Heterogeneous Systems: It Works!" in *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 6, no. 1. New York: Association for Computing Machinery, 2022, pp. 1–43. [Online]. Available: https://doi.org/10.1145/3508033 1, 2, 3, 4, 9, 10, 11, 19, 21, 25, 26

[4] M. van der Boor, S. C. Borst, J. S. H. van Leeuwaarden, and D. Mukherjee, "Scalable load balancing in networked systems: A survey of recent advances," 6 2018. [Online]. Available: http://arxiv.org/abs/1806.05444 3, 8

[5] M. Talagrand, *Mean Field Models for Spin Glasses*, 1st ed., ser. Ergebnisse der Mathematik und ihrer Grenzgebiete ; v. 54. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011. [Online]. Available: http://link.springer.com/10.1007/978-3-642-15202-3 3

[6] N. Gast, "Refinements of Mean Field Approximation," *Habilitation Thesis (HDR)*, 2020. [Online]. Available: http://polaris.imag.fr/nicolas.gast/pdfs/hdr_gast.pdf 3

[7] J. Abdul Jaleel, S. Doroudi, K. Gardner, and A. Wickeham, "A general "power-of-d" dispatching framework for heterogeneous systems," *Queueing Systems*, 2022. [Online]. Available: https://doi.org/10.1007/s11134-022-09736-z 3

[8] A. Mukhopadhyay and R. R. Mazumdar, "Analysis of Randomized Join-the-Shortest-Queue (JSQ) Schemes in Large Heterogeneous Processor-Sharing Systems," *IEEE Transactions on Control of Network Systems*, vol. 3, no. 2, pp. 116–126, 2016. 3

[9] K. Gardner, J. Abdul Jaleel, A. Wickeham, and S. Doroudi, "Scalable load balancing in the presence of heterogeneous servers," *Performance Evaluation*, vol. 145, p. 102151, 2021. [Online]. Available: https://doi.org/10.1016/j.peva.2020.102151 3

[10] N. Gast and B. Van Houdt, "A Refined Mean Field Approximation," in *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 1, 2017, p. 28. [Online]. Available: https://hal.inria.fr/hal-01622054 3

[11] N. Gast, "The power of two choices on graphs: The pair-approximation is accurate," *SIGMETRICS Performance Evaluation Review*, vol. 43, no. 2, pp. 69–71, 2015. 3, 6, 27, 32

[12] A. Budhiraja, D. Mukherjee, and R. Wu, "Supermarket model on graphs," *Annals of Applied Probability*, vol. 29, no. 3, pp. 1740–1777, 2019. 4, 8, 23, 28, 29, 30, 32, 42, 45

[13] D. Mukherjee, S. C. Borst, and J. S. van Leeuwaarden, "Asymptotically Optimal Load Balancing Topologies," in *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 2, no. 1. New York: Association for Computing Machinery, 2018, pp. 1–29. 4

[14] D. Rutten and D. Mukherjee, "Load Balancing Under Strict Compatibility Constraints," *Mathematics of Operations Research*, p. 31, 4 2022. [Online]. Available: http://pubsonline.informs.org/doi/10.1287/moor.2022.1258 4

[15] K. Ramanan, "Beyond mean-field limits for the analysis of large-scale networks," *Queueing Systems*, vol. 100, no. 3-4, pp. 345–347, 2022. [Online]. Available: https://doi.org/10.1007/s11134-022-09845-9 4, 35

[16] T. G. Kurtz, "Strong approximation theorems for density dependent Markov chains," *Stochastic Processes and their Applications*, vol. 6, no. 3, pp. 223–240, 1978. 7

[17] S. Allmeier and N. Gast, "Rmf Tool - A Library to Compute (Refined) Mean Field Approximation(s)," *SIGMETRICS Performance Evaluation Review*, vol. 49, no. 4, pp. 35–40, 2022. [Online]. Available: https://doi.org/10.1145/3543146.3543156 10, 20

[18] D. Hurtado-lange and S. T. Maguluri, "Throughput and delay optimality of power-of- d choices in inhomogeneous load balancing systems," *Operations Research Letters*, vol. 49, no. 4, pp. 616–622, 2021. [Online]. Available: https://doi.org/10.1016/j.orl.2021.06.010 13

[19] G. R. Grimmett and D. R. Stirzaker, *Probability and random processes*, 3rd ed. New York: Oxford university press, 2001. 13

[20] S. Foss and N. Chernova, "On the stability of a partially accessible multi-station queue with state-dependent routing ," *Queueing Systems*, vol. 29, pp. 55–73, 1998. [Online]. Available: https://doi.org/10.1023/A:1019175812444 23

[21] F. Cecchi, "Mean-Field Limits for Ultra-Dense Random-Access Networks," Ph.D. dissertation, Technische Universiteit Eindhoven., Eindhoven, 2018. 38

Appendix A

Proofs

A.1 Uniqueness

Proof of Lemma 1. We look at an arbitrary $(k,s) \in V_n \times \mathcal{S}$ element of (2.11). We have that $\phi^{(n)}(x,t) \in [0,1]^{n\times|\mathcal{S}|}$ and so if $f^{(n)}_{(k,s)}$ is a polynomial then it is bounded, continuous and Lipschitz continuous on $[0,1]$ by standard results. Then by the Cauchy–Lipschitz existence theorem we have that the solution of the initial value problem (2.11) must be unique on $[0,1]^{n\times|\mathcal{S}|}$. □

A.2 Independence of system size

This section contains the proofs of Theorems 3, 5, 8 and 9. All of these theorems make a similar statement along the idea that the mean field of a scaled system can be collapsed down into a simpler system of ODEs which is the mean field of the original pre-scaled system. There are subtleties to proving each one but they all follow the same basic plan:

Proof plan: We follow the following steps:

1. Consider an arbitrary scaled system. First show a form of exchangeability between the mean fields of like servers. The definition of similar servers depends on the specific context, but this usually emerges from the specific conditions imposed on the system. This is done in one of two ways:
    a) Show that the drifts are identical for two similar server (for example they both have the same neighbours). Then the mean field is the unique solution of (2.11) and so the ODEs for each server are identical up to an arbitrary relabelling of indices.
    b) If the drifts are not identical for two similar servers (for example they have different neighbours) then we expect that they still have the same structure drift (this is really what makes them similar type servers since the drift defines their growth). Thus if we assume the like servers have the same value at t_0 and then show that the drifts are equal at t_0, then since the drift is time homogeneous and the growth

of the servers is completely given by the drift, we have that they are equal for all $t \geq t_0$. If they have the same initial condition then they are equal for all time.

2. The exchangeability property means we can rewrite the system of ODEs (2.11) in terms of a representative for each group of like servers. The choice of representative is an arbitrary server in the group. This 'collapses' the system of ODEs (2.11) by removing the identical equations for like servers.

3. We show that the collapsed system of ODEs is independent of the scaling factor of the initial system. This is often tastily done in step 2. by rewriting the drift by group representatives the scaling factor is removed from the 'g' term in the drift which is often the only term that relates to the scaling.

4. Finally, notice that the original system of ODEs is 'equivalent' to the collapsed system whereby equivalent we mean any server in the scaled system has mean field defined by its representative. Then since the ODE is independent of the scaling factor, a scaled system of any degree could collapse into the same system of representatives.

□

A.2.1 Fixed proportions of service rates

Proof of Theorem 3. Let $X^{(aN)}$ be the proportionally scaled system of $X^{(N)}$ of arbitrary degree a. We first want to show that two servers of the same type have the same mean field i.e. $\phi^{(aN)}_{(k,s)}(x,t) = \phi^{(aN)}_{(k_1,s)}(x,t)$ for any $k, k_1 \in V_l^{(aN)}$ and for all $s \in \mathcal{S},\ t \in \mathbb{R}^+$ where ϕ is the mean field approximation and solution of (3.6).

Take two arbitrary servers $k, k' \in V_l^{(aN)}$, then we know $\mu_k = \mu^l = \mu_{k_1}$ and $g_s^{(n)}$ is independent of k and so we notice that $\phi^{(aN)}_{(k,s)}, \phi^{(aN)}_{(k_1,s)}$ are the unique solutions (Lemma 1) of the same differential equation but with an arbitrary relabelling of k to k_1. Since they have the same initial condition, they must be equal.

Now define the representative of group l, $\psi_{(l,s)} = \phi^{(aN)}_{(k,s)}$ for some arbitrary server $k \in V_l^{(aN)}$. We know that $\phi^{(aN)}_{(k,s)}$ is the unique solution of (3.6) and so we rewrite this in terms of $\psi_{(k,s)}(z,t)$ where z is the corresponding initial condition to x. Since all servers of the same type have the same initial condition, z is just the corresponding initial conditions for each of the m groups. We first rewrite $g_s^{(aN)}$ in terms of our new notation. We find for some $s \in \mathcal{S}$

$$g_s^{(aN)} = \sum_{k_1=1}^{aN} \sum_{s_1 \geq s} \frac{\phi^{(aN)}_{(k_1,s_1)}}{aN} = \sum_{s_1 \geq s} \left(\sum_{l=1}^{m} \left| V_l^{(aN)} \right| \frac{\psi_{(l,s_1)}}{aN} \right) = \sum_{s_1 \geq s} \left(\sum_{l=1}^{m} \alpha_l \psi_{(l,s_1)} \right).$$

Despite the suggestion from the notation, we actually see from the right hand side that $g_s^{(aN)}$ is independent of a and so we drop the dependence on aN (note that $m \leq aN$ but since $a \geq 1$ this condition is always satisfied thus being independent of N). For server k in group l then we write the drift as such

$$f^{(aN)}_{(k,s)}\left(\phi^{(aN)}\right) = \mu_k\phi^{(aN)}_{(k,s+1)} - \lambda\phi^{(aN)}_{(k,s)}(g^{(aN)}_{(k,s)} + g^{(aN)}_{(k,s+1)}) - \left(\mu_k\phi^{(aN)}_{(k,s)} - \lambda\phi^{(aN)}_{(k,s-1)}(g^{(aN)}_{(k,s-1)} + g^{(aN)}_{(k,s)})\right)\mathbb{1}_{\{s>0\}}. \tag{A.1}$$

$$= \mu^l(\psi_{(l,s+1)} - \psi_{(l,s)}\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_{(l,s)}(g_s + g_{s+1}) - \psi_{(l,s-1)}(g_{s-1} + g_s)\mathbb{1}_{\{s>0\}}\right). \tag{A.2}$$

For any server k in group l we have $\frac{d\phi^{(aN)}_{(k,s)}}{dt}(x,t) = \frac{d\psi_{(l,s)}}{dt}(z,t)$ and so the system of $aN \times b$ ODEs

$$\frac{d\phi^{(aN)}_{(k,s)}}{dt}(x,t) = f^{(aN)}_{(k,s)}\left(\phi^{(aN)}(x,t)\right)$$

collapses into the ODE system of representatives of $m \times b$ as such

$$\frac{d\psi_{(l,s)}}{dt} = \mu^l(\psi_{(l,s+1)} - \psi_{(l,s)}\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_{(l,s)}(g_s + g_{s+1}) - \psi_{(l,s-1)}(g_{s-1} + g_s)\mathbb{1}_{\{s>0\}}\right). \tag{A.3}$$

Since g_s is independent of a, the system (A.3) is also independent of a and thus the unique solution $\psi(z,t)$ is too. Since k was chosen arbitrary we can say that $\phi^{(aN)}_{(k,s)} = \psi_{(l,s)}$ for any $k \in V_l^{(aN)}$. Then since $\psi_{(l,s)}$ independent of a we can say that for two servers $k \in V_l^{(aN)}, k_1 \in V_l^{(N)}$ for all $s \in \mathcal{S}$ we have $\phi^{(aN)}_{(k,s)} = \psi_{(l,s)} = \phi^{(N)}_{(k_1,s)}$ for any $a \in \mathbb{N}$. Note that ψ is the unique solution of (A.3) by Lemma 1. □

A.2.2 Regular graphs

Proof of Theorem 5. We first want to show a form of 'exchangeability' by which the mean field approximation is equal for all servers i.e. $\phi^{(n)}_{(k,s)}(x,t) = \phi^{(n)}_{(k_1,s)}(x,t)$ for any $k, k_1 \in V_n$ and for all $s \in \mathcal{S}$ where ϕ is the new mean field approximation and therefore solution of (2.11).

In the proof of Theorem 3 found in A.2.1 we showed a similar exchangeability property and argued that the mean fields of each server were unique solutions to the same differential equation with an arbitrary relabelling of k, k_1. In the fully connected case this was easy to see since $g_s^{(n)}$ is independent of k. However, in this case on graphs $g^{(n)}_{(k,s)}$ depends on k. We instead use an inductive argument to show that we have exchangeability (2.b from the proof plan A.2). This relies on the fact that the drift is time homogeneous and only dependent on $\phi^{(n)}$ so if all $\phi^{(n)}_{(k,s)}(x,t)$ are equal, then the elements of the drift are equal. Therefore the growth of $\phi^{(n)}_{(k,s)}$ is entirely given by $\phi^{(n)}$ and so since the initial conditions are equal, they are equal for all time.

Assume that at some t_0 we have $\phi^{(n)}_{(k,s)}(x,t_0) = \phi^{(n)}_{(k',s)}(x,t_0)$ for all $k, k' \in V_n$ then by (2.11) we have

$$\frac{d\phi^{(n)}_{(k,s)}}{dt}(x,t_0) = f(\phi^{(n)}_{(k,s)}(x,t_0)) = f(\phi^{(n)}_{(k',s)}(x,t_0)) = \frac{d\phi^{(n)}_{(k',s)}}{dt}(x,t_0)$$

Therefore at t_0 when $\phi^{(n)}_{(k,s)}(x,t_0)$ coincides for all servers we see that their derivatives also coincide. Heuristically, if we take some infinitesimal step in t which we call t_1 then we see that once again all $\phi^{(n)}_{(k,s)}(x,t_1) = \phi^{(n)}_{(k',s)}(x,t_1)$ for all $k,k' \in V$ since their derivative was equal. Now since they are equal again, we know their derivative is equal at t_0. Thus inductively we can see that for all $t > t_0$ we have $\phi^{(n)}_{(k,\cdot)}(x,t) = \phi^{(n)}_{(k',\cdot)}(x,t)$ for all $k,k' \in V$. Finally we note that due to the initial condition we know that at $t = 0$ we have $\phi^{(n)}_{(k,s)}(x,0) = \phi^{(n)}_{(k',s)}(x,0)$ for all $k,k' \in V$ and so we see that we have said desired 'exchangeability'.

Fix $k \in V_n$ arbitrarily, we now define the representative $\psi_s(z,t) = \phi^{(n)}_{(k,s)}(x,t)$ where z is the corresponding initial condition for the k-th server i.e the vector x_k. We know that $\phi^{(n)}$ is the solution of the differential equation (2.11), where $f^{(n)}$ is the drift. We wish to explicitly write out ODE (2.11) as a system in $\phi^{(n)}_{(k,s)}$ and then rewrite this in terms of $\psi_s(z,t)$. Looking at $g^{(n)}_{(k,s)}$ we get

$$g^{(n)}_{(k,s)} = \sum_{k_1 \in \mathcal{N}^{(n)}_k} \sum_{s_1 \geq s} \frac{\phi^{(n)}_{(k_1,s_1)}}{\left|\mathcal{N}^{(n)}_k\right|} = \sum_{k_1 \in \mathcal{N}^{(n)}_k} \sum_{s_1 \geq s} \frac{\phi^{(n)}_{(k,s_1)}}{\left|\mathcal{N}^{(n)}_k\right|} = \sum_{k_1 \in \mathcal{N}^{(n)}_k} \sum_{s_1 \geq s} \frac{\psi_{s_1}}{\left|\mathcal{N}^{(n)}_k\right|} = \sum_{s_1 \geq s} \psi_{s_1} := \tilde{g}_s$$

by the above exchangeability. We see that despite the notation, this is independent of n and since k was fixed arbitrarily, it is also independent of the choice of k. Using this we then write the drift into the form

$$f^{(n)}_{(k,s)}(\phi^{(n)}) = \mu\phi^{(n)}_{(k,s+1)} - \lambda\phi^{(n)}_{(k,s)}(g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}) - \left(\mu_k\phi^{(n)}_{(k,s)} - \lambda\phi^{(n)}_{(k,s-1)}(g^{(n)}_{(k,s-1)} + g^{(n)}_{(k,s)})\right)\mathbb{1}_{\{s>0\}}. \tag{A.4}$$

$$= \mu(\psi_{s+1} - \psi_s\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_s(\tilde{g}_s + \tilde{g}_{s+1}) - \psi_{s-1}(\tilde{g}_{s-1} + \tilde{g}_s)\mathbb{1}_{\{s>0\}}\right) \tag{A.5}$$

Since k was chosen arbitrarily and all $\phi^{(n)}_{(k,\cdot)}$ are equal by exchangeability, we have that any $k \in V_n$ $\phi^{(n)}_{(k,s)}(x,t) = \psi_s(z,t)$ and hence also $\frac{d\phi^{(n)}_{(k,s)}}{dt}(x,t) = \frac{d\psi_s}{dt}(z,t)$ which combined with the above drift gives that the differential equation $\frac{d\phi^{(n)}_{(k,s)}}{dt}(x,t) = f^{(n)}_{(k,s}(\phi^{(n)}(x,t))$ is equivalent to

$$\frac{d\psi_s}{dt} = \mu(\psi_{s+1} - \psi_s\mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_s(\tilde{g}_s + \tilde{g}_{s+1}) - \psi_{s-1}(\tilde{g}_{s-1} + \tilde{g}_s)\mathbb{1}_{\{s>0\}}\right). \tag{A.6}$$

Finally notice that the solution of the system of ODEs (A.6) is unique by Lemma 1. Thus $\phi^{(n)}_{(k,s)}(x,t) = \psi_s(z,t)$ for all $k \in V_n$ where $\psi(z,t)$ the solution of the system of ODEs (A.6) as required. □

A.2.3 Regular graphs with heterogeneous service rates

Proof of Theorem 8. Our approach is to show that the mean field for system $X^{(N)}$ with regular graph structure G_N of degree d and m types of servers satisfies (5.6) and is therefore 'independent' of the size of the system n and degree of the graph d. This process is nearly identical to the proof of Theorems 3 and 5 found in A.2.1 and A.2.2 but more generalised and so less clear. We follow the same proof plan A.2.

We first show 'exchangeability' between servers of the same type. For this we use the argument 2.b in proof plan A.2 which was written in more detail in the proof of theorem 5 in A.2.2. Servers of the same type do not necessarily have the same connections and so it is not a case of an arbitrary relabelling of indices. However, the definition of type is very strong and by construction means the drifts for two servers of the same type have the same value if all servers of the same type start from the same initial conditions.

Briefly the argument is as follows. If at some t_0 all servers of the same type have the same mean field, then clearly for two servers of the same type k, k' their 'g's are equal i.e. $g^{(n)}_{(k,\cdot)} = g^{(n)}_{(k',\cdot)}|_{t=0}$ since by construction they both have the same type neighbours. Then since they are of the same type they have the same processing rate and so their drifts are equal. The drift is equal to the derivative of their mean field and so they have the same growth at time $t = 0$, thus they will still be equal after an infinitesimal time step t_1 which repeats for all time inductively. We know at $t = 0$ all servers of the same type have the same initial condition and thus the exchangeability holds.

Choose an arbitrary $k \in V^{(n)}_l$ define type representative $\psi_{(l,\cdot)} = \phi^{(n)}_{(k,\cdot)}$. We now rewrite the ODE (2.11) for k in terms of ψ. We begin by looking at the $g^{(n)}_{(k,s)}$. We use exchangeability of servers of the same type to change the summation from the neighbourhood of servers $\mathcal{N}^{(n)}_k$ to the neighbourhood of types $\mathcal{L}_l$ which we recall is the set of all types of servers in the neighbourhood of type l servers. For a type l server $k \in V^{(n)}_l$ we have that

$$g^{(n)}_{(k,s)} = \sum_{k_1 \in \mathcal{N}^{(n)}_k} \sum_{s_1 \geq s} \frac{\phi^{(n)}_{(k_1,s_1)}}{d} = \sum_{j \in \mathcal{L}_l} \sum_{s_1 \geq s} \frac{\alpha^j_l d \psi_{(j,s_1)}}{d} = \sum_{j \in \mathcal{L}_l} \sum_{s_1 \geq s} \alpha^j_l \psi_{(j,s_1)} := \tilde{g}_{(l,s)}. \tag{A.7}$$

where α^j_l was defined as the proportion of type j servers in the neighbourhoods of type l servers. Also note that all the neighbourhood sizes are that of the degree of the graph which gives the divisor d. We define $\tilde{g}_{(l,s)}$ in (A.7) and notice that this is independent of the system size n and degree of the graph d. We now rewrite the (k, s) element of the drift for $k \in V^{(n)}_l$ which gives

$$f^{(n)}_{(k,s)}\left(\phi^{(n)}\right) = \mu_k \phi^{(n)}_{(k,s+1)} - \lambda \phi^{(n)}_{(k,s)} (g^{(n)}_{(k,s)} + g^{(n)}_{(k,s+1)}) - \left(\mu_k \phi^{(n)}_{(k,s)} - \lambda \phi^{(n)}_{(k,s-1)} (g^{(n)}_{(k,s-1)} + g^{(n)}_{(k,s)})\right) \mathbb{1}_{\{s>0\}} \tag{A.8}$$

$$= \mu^l (\psi_{(l,s+1)} - \psi_{(l,s)} \mathbb{1}_{\{s>0\}}) - \lambda \left(\psi_{(l,s)} (\tilde{g}_{(l,s)} + \tilde{g}_{(l,s+1)}) - \psi_{(l,s-1)} (\tilde{g}_{(l,s-1)} + \tilde{g}_{(l,s)}) \mathbb{1}_{\{s>0\}}\right). \tag{A.9}$$

Since k was chosen arbitrarily and we have exchangeability between servers of the same type, we have for all $k \in V^{(n)}_l$ $\phi^{(n)}_{(k,\cdot)}(x,t) = \psi_{(l,\cdot)}(z,t)$. Therefore we also have $\frac{d\phi^{(n)}_{(k,s)}}{dt}(x,t) = \frac{d\psi_{(l,s)}}{dt}(z,t)$ and so the $n \times b$ system of ODEs

$$\frac{d\phi^{(n)}_{(k,s)}}{dt}(x,t) = f^{(n)}_{(k,s)}\left(\phi^{(n)}(x,t)\right)$$

is equivalent to the $m \times b$ system of representatives

$$\frac{d\psi_{(l,s)}}{dt} = \mu^l (\psi_{(l,s+1)} - \psi_{(l,s)} \mathbb{1}_{\{s>0\}}) - \lambda \left(\psi_{(l,s)} (\tilde{g}_{(l,s)} + \tilde{g}_{(l,s+1)}) - \psi_{(l,s-1)} (\tilde{g}_{(l,s-1)} + \tilde{g}_{(l,s)}) \mathbb{1}_{\{s>0\}}\right). \tag{A.10}$$

Since $\tilde{g}_{(l,s)}$ is independent of n, d, the system (A.10) is also independent of n, d and thus the unique solution $\psi(z,t)$ is too. Since k was chosen arbitrary we can say that $\phi^{(n)}_{(k,s)} = \psi_{(l,s)}$ for any $k \in V_l^{(n)}$.

□

A.2.4 Cluster graphs

Proof of Theorem 9. Our approach here is to study the cluster scaled system $X^{(aN)}$ and then show that for servers in the same cluster, their mean field is given by the mean field of the initial system. We begin just like proof of Theorem 3 and the proof plan A.2 by showing 'exchangeability' within the same cluster. This means we must show that two servers in the same cluster have the same mean field i.e. $\phi^{(aN)}_{(k_1,\cdot)}(x,t) = \phi^{(aN)}_{(k_2,\cdot)}(x,t)$ for any $k_1, k_2 \in V_k^{(aN)}$ and for all $t \in \mathbb{R}^+$.

Take two arbitrary servers $k_1, k_2 \in V_k^{(aN)}$ from cluster k, then we know they have the same service rate. Furthermore they are connected to exactly the same servers by definition of the cluster graph and so $g_{(k_1,\cdot)} = g_{(k_2,\cdot)}$. Therefore $\phi^{(aN)}_{(k_1,s}$ and $\phi^{(aN)}_{(k_2,s)}$ are the unique solutions (Lemma 1) to the same ODE up to an arbitrary relabelling of k_1, k_2 and so must be equal.

We now want to use exchangeability to show that the (k,s) element of the ODE (2.11) is independent of a. As shown in proof of Theorem 3 found in(A.2.1), this equates to showing that $g^{(aN)}_{(k,s)}$ is actually independent of a and finding a new system of ODEs which collapses the exchangeable servers into one ODE for each cluster. In the neighbourhood of a server in the k-th cluster we have a many servers of each type of connected cluster i.e. we have the set of connected clusters $\mathcal{N}_k^{(N)}$ and then a many servers for each cluster. This means that

$$\left|\mathcal{N}_k^{(aN)}\right| = a\left|\mathcal{N}_k^{(N)}\right|.$$

Let k_1 represent an arbitrary server in the k-th cluster, then

$$\begin{aligned}
g^{(aN)}_{(k_1,s)} &= \frac{1}{2}\sum_{k'\in\mathcal{N}_{k_1}^{(aN)}}\sum_{s'\geq s}\phi^{(aN)}_{(k',s')}\left(\frac{1}{\left|\mathcal{N}_{k_1}^{(aN)}\right|}+\frac{1}{\left|\mathcal{N}_{k'}^{(n)}\right|}\right) && \text{(A.11)}\\
&= \frac{1}{2}\sum_{k'\in\mathcal{N}_{k_1}^{(N)}}\sum_{s'\geq s}a\phi^{(aN)}_{(k',s')}\left(\frac{1}{a\left|\mathcal{N}_{k_1}^{(N)}\right|}+\frac{1}{a\left|\mathcal{N}_{k'}^{(N)}\right|}\right) && \text{(A.12)}\\
&= \frac{1}{2}\sum_{k'\in\mathcal{N}_{k_1}^{(N)}}\sum_{s'\geq s}\phi^{(aN)}_{(k',s')}\left(\frac{1}{\left|\mathcal{N}_{k_1}^{(N)}\right|}+\frac{1}{\left|\mathcal{N}_{k'}^{(N)}\right|}\right) && \text{(A.13)}\\
&= g^{(N)}_{(k_1,s)} = g^{(N)}_{(k,s)}. && \text{(A.14)}
\end{aligned}$$

This shows that g in the cluster scaled system of degree a is equal to g in the initial system. Notice that there is only one server in the k-th cluster of the initial system, i.e. the k-th server which justifies line (A.14). Much like in the proof of Theorem 3, we can now rewrite the drift and so we find a many replications of the same ODE for servers in the same cluster. If we write the system of representatives, one for each cluster then we get the system of ODEs that describes the mean field $\phi^{(N)}$, of the original N sized system (with unique solution Lemma 1). This follows from the fact that $g^{(aN)}_{(k_1,s)} = g^{(N)}_{(k,s)}$ for $k_1 \in V_k^{(aN)}$. Therefore the mean field of a server in cluster k of the system $X^{(aN)}$ is equal to server k of the pre-scaled system $X^{(N}$. □

A.3 Limit of expected occupancy state

A.3.1 Fixed proportions of service rates

Proof of Theorem 4. Using Definition 5 we look at the expectation of the occupation state for a server of type l and length s in the a-th system of the sequence $(X^{(d(a)N)})_{a\in\mathbb{N}}$. Using linearity of the expectation we have

$$\mathbb{E}\left[x^{(d(a)N)}_{(l,s)}(t)\right] = \mathbb{E}\left[\frac{1}{\alpha_l d(a)N}\sum_{k\in V_l^{(d(a)N)}} X^{(d(a)N)}_{(k,s)}\right] = \frac{1}{\alpha_l d(a)N}\sum_{k\in V_l^{(d(a)N)}} \mathbb{E}\left[X^{(d(a)N)}_{(k,s)}\right].$$

We can now investigate $\mathbb{E}\left[X^{(d(a)N)}_{(k,s)}\right]$ using Theorem 2 to obtain

$$\frac{1}{\alpha_l d(a)N}\sum_{k\in V_l^{(d(a)N)}} \mathbb{E}\left[X^{(d(a)N)}_{(k,s)}\right] = \frac{1}{\alpha_l d(a)N}\sum_{k\in V_l^{(d(a)N)}} \phi^{(d(a)N)}_{(k,s)} + O\left(\frac{1}{d(a)N}\right),$$

where $\phi^{(d(a)N)}$ is the mean field for the a-th system in the sequence of proportionally scaled systems. Now by Theorem 3, for any $a \in \mathbb{N}$ we have $\phi^{(d(a)N)}_{(k,s)} = \psi_{(l,s)}$ for all $k \in V_l^{(d(a)N)}$ and $s \in \mathcal{S}$. Since we sum over exactly $k \in V_l^{(d(a)N)}$, we substitute ψ into our summation which gives

$$\begin{aligned}\frac{1}{\alpha_l d(a)N}\sum_{k\in V_l^{(d(a)N)}} \psi_{(l,s)} + O\left(\frac{1}{d(a)N}\right) &= \frac{1}{\alpha_l d(a)N}\left|V_l^{(d(a)N)}\right|\left[\psi_{(l,s)} + O\left(\frac{1}{d(a)N}\right)\right]\\ &= \psi_{(l,s)} + O\left(\frac{1}{d(a)N}\right).\end{aligned}$$

Resulting in the expectation of the a-th occupation state given by

$$\mathbb{E}\left[x^{(d(a)N)}_{(l,s)}(t)\right] = \psi_{(l,s)} + O\left(\frac{1}{d(a)N}\right).$$

Then as $a \to \infty$, $d(a) \to \infty$ and thus

$$\lim_{a\to\infty}\mathbb{E}\left[x^{(d(a)N)}_{(l,s)}(t)\right] = \psi_{(l,s)}.$$

Finally recall that as given in Theorem 3, $\psi_{(l,s)}$ is the unique solution of the ODE (3.7).

□

A.3.2 Regular graphs

Proof of Theorem 6. This proof is nearly identical to the proof of Theorem 4 found in A.3.1.

Using definition 6 we look at the expectation of the occupation state for a server of queue length s in the a-th system of the sequence $(X^{(n(a))})_{a\in\mathbb{N}}$. Recall that the a-th system is of $n(a)$ servers and exists on a graph $G_{n(a)} = (V_{n(a)}, E_{n(a)})$ of degree $d(a)$ such that both $n(a), d(a) \to \infty$ as $a \to \infty$. Also note that $n(a) \geq d(a)$ for all a since the degree cannot be higher than the number of servers. Using linearity of the expectation we have

$$\mathbb{E}\left[x_s^{(n(a))}(t)\right] = \mathbb{E}\left[\frac{1}{n(a)} \sum_{k\in V_{(n(a))}} X_s^{(n(a))}\right] = \frac{1}{n(a)} \sum_{k\in V_{(n(a))}} \mathbb{E}\left[X_s^{(n(a))}\right].$$

We can now investigate $\mathbb{E}\left[X_s^{(n(a))}\right]$ using Proposition 1 to obtain

$$\frac{1}{n(a)} \sum_{k\in V_{(n(a))}} \mathbb{E}\left[X_s^{(n(a))}\right] = \frac{1}{n(a)} \sum_{k\in V_{(n(a))}} \phi_{(k,s)}^{(n(a))} + O\left(\frac{1}{d(a)}\right),$$

where $\phi^{(n(a))}$ is the mean field for the a-th system in the sequence of proportionally scaled systems. Now by Theorem 5, for any $a \in \mathbb{N}$ we have $\phi_{(k,s)}^{(n(a))} = \psi_s$ for all $k \in V^{(n(a))}$ and $s \in \mathcal{S}$ and ψ solution of (4.6). Substituting ψ into our summation gives

$$\begin{aligned}\frac{1}{n(a)} \sum_{k\in V_{(n(a))}} \psi_s + O\left(\frac{1}{d(a)}\right) &= \frac{1}{n(a)} \left|V_{(n(a))}\right| \left[\psi_s + O\left(\frac{1}{d(a)}\right)\right] \\ &= \psi_s + O\left(\frac{1}{d(a)}\right).\end{aligned}$$

Resulting in the expectation of the a-th occupation state given by

$$\mathbb{E}\left[x_{(s}^{(n(a))}(t)\right] = \psi_s + O\left(\frac{1}{d(a)}\right).$$

Then as $a \to \infty$, $d(a) \to \infty$ and thus

$$\lim_{a\to\infty} \mathbb{E}\left[x_s^{(n(a))}(t)\right] = \psi_s.$$

Finally recall that as given in Theorem 3, ψ_s is the unique solution of the ODE (3.7).

□

A.3.3 Cluster graphs

Proof of Theorem 10. We look at the expectation of the occupation state (Definition 9) for cluster k and queue length s in the a-th system of the sequence $(X^{(d(a)N)})_{a\in\mathbb{N}}$. Using linearity of the expectation we have

$$\mathbb{E}\left[x_{(k,s)}^{(d(a)N)}(t)\right] = \mathbb{E}\left[\frac{1}{d(a)} \sum_{k_1\in V_k^{(d(a)N)}} X_{(k_1,s)}^{(d(a)N)}\right] = \frac{1}{d(a)} \sum_{k_1\in V_k^{(d(a)N)}} \mathbb{E}\left[X_{(k_1,s)}^{(d(a)N)}\right].$$

We can now analyse $\mathbb{E}\left[X_{(k,s)}^{(d(a)N)}\right]$ using Proposition 1. Let $\mathcal{N}_k^{(N)}$ be the neighbourhood of server k in the original system thus the neighbourhood size for a server k_1 in cluster k of the $a-$th system is equal to $d(a)|\mathcal{N}_k^{(N)}|$. Since $d(a) \to \infty$ as $a \to \infty$ we see the neighbourhoods grow to infinity and therefore we can apply Proposition 1. Hence we obtain

$$\frac{1}{d(a)} \sum_{k_1 \in V_k^{(d(a)N)}} \mathbb{E}\left[X_{(k_1,s)}^{(d(a)N)}\right] = \frac{1}{d(a)} \sum_{k_1 \in V_k^{(d(a)N)}} \phi_{(k,s)}^{(d(a)N)} + O\left(\frac{1}{d(a)|\mathcal{N}_k^{(N)}|}\right),$$

where $\phi^{(d(a)N)}$ is the mean field for the a-th system in the sequence of proportionally scaled systems. Now by Theorem 9, for any $a \in \mathbb{N}$ we have $\phi_{(k_1,s)}^{(d(a)N)} = \phi_{(k,s)}^{(N)}$ for all $k_1 \in V_k^{(d(a)N)}$ and $s \in \mathcal{S}$. Since we sum over exactly $k \in V_k^{(d(a)N)}$, we substitute $\phi^{(N)}$ into our summation which gives

$$\begin{aligned}\frac{1}{d(a)} \sum_{k_1 \in V_k^{(d(a)N)}} \phi_{(k,s}^{(N)} + O\left(\frac{1}{d(a)|\mathcal{N}_k^{(N)}|}\right) &= \frac{1}{d(a)} \left|V_k^{(d(a)N)}\right| \left[\phi_{(k,s)}^{(N)} + O\left(\frac{1}{d(a)|\mathcal{N}_k^{(N)}|}\right)\right] \\ &= \phi_{(k,s)}^{(N)} + O\left(\frac{1}{d(a)|\mathcal{N}_k^{(N)}|}\right).\end{aligned}$$

Resulting in the expectation of the a-th occupation state given by

$$\mathbb{E}\left[x_{(l,s)}^{(d(a)N)}(t)\right] = \phi_{(k,s)}^{(N)} + O\left(\frac{1}{d(a)|\mathcal{N}_k^{(N)}|}\right).$$

Then as $a \to \infty$, $d(a) \to \infty$ and thus

$$\lim_{a\to\infty} \mathbb{E}\left[x_{(k,s)}^{(d(a)N)}(t)\right] = \phi_{(k,s)}^{(N)}.$$

where $\phi^{(N)}$ is the new mean field of the original system $X^{(N)}$ $\square$

A.4 Conditions on graphs

Proof of Theorem 7. The backwards direction follows immediately from Theorem 5. In the forwards direction we start by assuming that the condition $\phi_{(k,s)}^{(n)} = \psi_s$ for all $k \in V_n$. Recall that $\phi^{(n)}$ is the new mean field of the homogeneous system on a graph G_n and ψ_s is the mean field for the same system but on the fully connected graph. Let the initial condition of $\phi^{(n)}$ be given by $z = (x, ..., x)$. Since the two mean fields are equal for all t we find their derivatives with respect to t must also be equal which gives

$$\frac{d\phi_{(k,s)}^{(n)}}{dt}(z,t) = \frac{d\tilde{\psi}_s}{dt}(x,t) \implies f_{(k,s)}^{(n)}(\phi^{(n)}(z,t)) = \frac{d\psi_s}{dt}(x,t) \tag{A.15}$$

We know that the sizes of the neighbourhoods are captured within the $g_{(k,s)}$ term of the drift on the left hand side. We want to try and extract this. The set of graphs that satisfy (A.15) (which we call $\mathcal{G}$) must be a subset of those that satisfy it when the initial condition starts from the empty state. Let $\psi_0(x,0) = 1, \psi_s(x,0) = 0$ for all $s \geq 1$ and then for all $k \in V_n$ let $\phi^{(n)}_{(k,0)}(z,0) = 1$ and $\phi^{(n)}_{(k,s)}(z,0) = 0$ for all $s \geq 1$. Then we first find $\frac{d\psi_s}{dt}(x,0)$ as such

$$\frac{d\psi_s}{dt}(x,t) = \mu(\psi_{s+1} - \psi_s \mathbb{1}_{\{s>0\}}) - \lambda\left(\psi_s(\tilde{g}_s + \tilde{g}_{s+1}) - \psi_{s-1}(\tilde{g}_{s-1} + \tilde{g}_s)\mathbb{1}_{\{s>0\}}\right) \tag{A.16}$$

$$\frac{d\psi_0}{dt}(x,0) = \mu(-1 \cdot \mathbb{1}_{\{s>0\}}) - \lambda\left(1 \cdot (1+0)\right) = -\lambda \tag{A.17}$$

$$\frac{d\psi_s}{dt}(x,0) = 0 \quad \text{for } s \geq 1 \tag{A.18}$$

which we obtained by observing that $\tilde{g}_0 = 1$ and $\tilde{g}_s = 0$ for $s \geq 1$. Now we look at the $(k,0)$-th element of the drift for $\phi^{(n)}$,

$$f^{(n)}_{(k,0)}(\phi^{(n)}(x,t)) = \mu\phi^{(n)}_{(k,1)} - \lambda\phi^{(n)}_{(k,0)}(g^{(n)}_{(k,0)} + g^{(n)}_{(k,1)}) \tag{A.19}$$

then at $t = 0$ we find

$$f^{(n)}_{(k,0)}(\phi^{(n)}(x,0)) = \mu\phi^{(n)}_{(k,1)} - \lambda\phi^{(n)}_{(k,0)}(g^{(n)}_{(k,0)}(t=0) + g^{(n)}_{(k,1)}(t=0)) \tag{A.20}$$

$$= -\lambda(g^{(n)}_{(k,0)}(t=0) + g^{(n)}_{(k,1)}(t=0)). \tag{A.21}$$

Now studying $g^{(n)}_{(k,0)}, g^{(n)}_{(k,1)}$ gives

$$g^{(n)}_{(k,0)} = \frac{1}{2}\sum_{k_1 \in \mathcal{N}_k}\sum_{s_1=0}^{b}\phi^{(n)}_{(k,s_1)}(x,0)\left(\frac{1}{|\mathcal{N}_k|} + \frac{1}{|\mathcal{N}_{k_1}|}\right) \tag{A.22}$$

$$= \frac{1}{2}\sum_{k_1 \in \mathcal{N}_k}\left(\frac{1}{|\mathcal{N}_k|} + \frac{1}{|\mathcal{N}_{k_1}|}\right) \tag{A.23}$$

$$g^{(n)}_{(k,1)} = \frac{1}{2}\sum_{k_1 \in \mathcal{N}_k}\sum_{s_1=1}^{b}\phi^{(n)}_{(k,s_1)}(x,0)\left(\frac{1}{|\mathcal{N}_k|} + \frac{1}{|\mathcal{N}_{k_1}|}\right) \tag{A.24}$$

$$= 0. \tag{A.25}$$

Combining this gives means that we can rewrite (A.15) at time $t = 0$ from the empty system into the form

$$f^{(n)}_{(k,s)}(\phi^{(n)}(z,0)) = \frac{d\psi_s}{dt}(x,0) \tag{A.26}$$

$$\implies \frac{-\lambda}{2}\sum_{k_1 \in \mathcal{N}_k}\left(\frac{1}{|\mathcal{N}_k|} + \frac{1}{|\mathcal{N}_{k_1}|}\right) = -\lambda \tag{A.27}$$

$$\implies \frac{1}{2}\sum_{k_1 \in \mathcal{N}_k}\left(\frac{1}{|\mathcal{N}_k|} + \frac{1}{|\mathcal{N}_{k_1}|}\right) = 1 \tag{A.28}$$

$$\implies \frac{1}{2} + \frac{1}{2}\sum_{k_1 \in \mathcal{N}_k}\frac{1}{|\mathcal{N}_{k_1}|} = 1 \tag{A.29}$$

$$\implies \sum_{k_1 \in \mathcal{N}_k}\frac{1}{|\mathcal{N}_{k_1}|} = 1 \qquad \text{for all } k \in V. \tag{A.30}$$

We now need to find all graphs that satisfy (A.30) for which we know our desired set of graphs $\mathcal{G}$ must be a subset of. We note a few properties of this system of equations.

i) The number of terms in the k-th equation is the size of the neighbourhood, $|\mathcal{N}_k|$

ii) If we find a neighbour has neighbourhood size N then we know we have another equation centred around this point with N many terms

We wish to prove the following statement: *The only solutions to* (A.30) *are disjoint unions of regular graphs i.e. for any server* k, $|\mathcal{N}_k^{(n)}| = |\mathcal{N}_{k'}^{(n)}|$ *for all* $k' \in \mathcal{N}_k$.

We provide the proof for the JSQ(2) without replacement policy since this is leading. However we provide a note on how to adjust it to hold for the case with replacement too. We prove the statement inductively.

JSQ without replacement
In this case each server is not in its own interaction neighbourhood. We first look at the base case where for some $k \in V_n$, $|\mathcal{N}_k| = 1$. Therefore there is one server in its neighbourhood which we call k_1, thus (A.30) becomes $1 = 1/|\mathcal{N}_{k_1}|$ which implies $|\mathcal{N}_{k_1}| = 1$. Therefore if there are any servers with a neighbourhood of size one, then it must also only be connected to another server with neighbourhood size one (which forms a disjoint regular subgraph of just the two connected servers).

We now assume that for all neighbourhood sizes $|\mathcal{N}_k^{(n)}| \leq m$ for some $m \in \mathbb{N}$, all the neighbours of k have the same neighbourhood size as k i.e. for $k_i \in \mathcal{N}_k^{(n)}, |\mathcal{N}_{k_i}^{(n)}| = |\mathcal{N}_k^{(n)}|$. Now we take the inductive step, suppose $|\mathcal{N}_k^{(n)}| = m + 1$ then it cannot have any neighbours with size $m' = |\mathcal{N}_{k_i}^{(n)}| \leq m$ since all the neighbours of k_i must then also have neighbourhood size m' which is a contradiction since $k \in \mathcal{N}_{k_i}^{(n)}$. Therefore it must have neighbours with larger neighbourhoods. Let the $m+1$ neighbours of server k be given by $k_1, ..., k_{m+1}$ Suppose $|\mathcal{N}_{k_{m+1}}^{(n)}| > m + 1$ then from (A.30) we have

$$1 = \frac{1}{|\mathcal{N}_{k_1}^{(n)}|} + ... + \frac{1}{|\mathcal{N}_{k_{m+1}}^{(n)}|} < \frac{1}{|\mathcal{N}_{k_1}^{(n)}|} + ... + \frac{1}{|\mathcal{N}_{k_m}^{(n)}|} + \frac{1}{m+1}$$
$$\implies \frac{m}{m+1} < \frac{1}{|\mathcal{N}_{k_1}^{(n)}|} + ... + \frac{1}{|\mathcal{N}_{k_m}^{(n)}|} \leq \max\left(\frac{1}{|\mathcal{N}_{k_1}^{(n)}|} + ... + \frac{1}{|\mathcal{N}_{k_m}^{(n)}|}\right) = \frac{m}{m+1}$$

this gives a contraction and so $|\mathcal{N}_{k_{m+1}}^{(n)}| = m+1$. Repeating this gives that all neighbours have neighbourhood size $m + 1$ and so the inductive step holds. Due to the base step for $m = 1$ and inductive step, we find this holds for all $m \in \mathbb{N}$ by mathematical induction.

Therefore for any server with any neighbourhood size, all of their neighbours must have the same neighbourhood size which means we have a disjoint regular subgraph of some larger graph G_n. Hence the only solutions of (A.30) are disjoint unions of regular graphs. Furthermore we know the set of all graphs satisfying (A.15) $\mathcal{G}$ must be a subset of the set of disjoint unions of regular graphs. Thus if a graph belongs to $\mathcal{G}$ it must be (a disjoint union of) regular graph(s).

JSQ(2) with replacement
The difference with this policy is that a server k belongs to its own neighbourhood. this changes the base case to mean when the neighbourhood size is 1 we have a solitary server

and when the neighbourhood size is 2 we have a subgraph of two connected servers. The reasoning in the inductive step holds, only the k_{m+1} server is in fact k which we know has size $m + 1$ but the idea is to obtain the same contradiction with the maximum. The proceeding reasoning is the same. □

Appendix B

Derivations

B.1 Power-of-d

Derivation of (3.9). Consider a system of size n according to Definition 4. Let $Q_{(l,s)}^{(n)}$ be the number of servers of type l that have queue length s or more. Then we have state space

$$q = \left(q_{(l,s)}\right)_{1\leq l\leq m, s\in\mathcal{S}} \in \left\{\left(Q_{(l,s)}/n\right)_{1\leq l\leq m, s\in\mathcal{S}} \mid Q_{(l,s)} \in \{0, ..., n\}\right\} \tag{B.1}$$

Let $e_{(l,s)}$ be the same sized matrix with a 1 in the element (l, s) and zero elsewhere. Then we find the following transition rates

$$q \to q + e_{(l,s-1)}/n - e_{(l,s)}/n \quad \text{at rate} \quad n\mu^l \left(q_{(l,s)} - q_{(l,s+1)}\right) \tag{B.2}$$
$$q \to q + e_{(l,s+1)}/n - e_{(l,s)}/n \quad \text{at rate} \quad n\lambda P_s \tag{B.3}$$

where P_s is given by

$$P_{(l,s)} = \sum_{i=1}^{d}\sum_{j=1}^{d-i} \frac{i}{i+j} B(d,j,i) p_{s+1}^{d-i-j} (q_{(l,s)} - q_{(l,s+1)})^i ((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))^j \tag{B.4}$$

$$p_s = \sum_{l=1}^{m} q_{(l,s)} \tag{B.5}$$

$$B(d,j,i) = \frac{d!}{i!j!(d-i-j)!} \tag{B.6}$$

We now justify these transition rates. The rate (B.2) describes a customer leaving a type l server with queue length s henceforth an (l, s) server. We have $n(q_{(l,s)} - q_{(l,s+1)})$ many type l servers of queue length s and so a departure occurs at rate $\mu^l \cdot n(q_{(l,s)} - q_{(l,s+1)})$. The rates (B.3) require a few steps to explain.

We first want to compute $P_{(l,s)}$ which is the probability that an arriving job is assigned to a type (l, s) server. This is the sum over the probability of sampling the specific d servers multiplied by the probability of an arrival to a type (l, s) server given the d servers that were sampled. Clearly this is zero if we sample a server with queue length shorter than s and 1 we sample a type (l, s) server and only other servers that are longer. The difficulty arises with the tie breaks which are uniform, therefore we must know how many type (l, s) and type (l', s) servers $l' \neq l$ we have. Therefore we consider three different categories of servers. When a job arrive we sample d many servers of which we can choose servers with queue lengths longer than s, type (l, s) servers or type (l', s) servers $l' \neq l$. Since we have replacement, we just consider the probability to sample each of these type servers once and then take the product of probabilities. The probability to sample servers with queue lengths longer than s is p_{s+1}. To sample a type (l, s) is $(q_{(l,s)} - q_{(l,s+1)})$ and type (l', s), $l' \neq l$ $((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))$. Then consider the number of ways we can sample these categories which is the number of ways to arrange d objects of which $i, j, d - i - j$ are the same which gives (B.6). The probability to sample $i, j, d - i - j$ of our respective objects and for the arrival to be assigned to a type (l, s) is given by

$$\frac{i}{i+j} B(d, j, i) p_{s+1}^{d-i-j} (q_{(l,s)} - q_{(l,s+1)})^i ((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))^j.$$

The $i/(i + j)$ is the probability that a tie break is won by a type (l, s) server. If $i = 0$ this is equal to zero and so we sum over $i = 1, ..., d$ and $j = 0, ..., d - i$ to obtain $P_{(l,s)}$

$$P_{(l,s)} = \sum_{i=1}^{d} \sum_{j=1}^{d-i} \frac{i}{i+j} B(d, j, i) p_{s+1}^{d-i-j} (q_{(l,s)} - q_{(l,s+1)})^i ((p_s - p_{s+1}) - (q_{(l,s)} - q_{(l,s+1)}))^j$$

Now by the splitting of the Poisson process the arrival rate to type (l, s) servers if given by $n\lambda P_{(l,s)}$ which gives (B.3).

Much like in the population process approach in Section 2.3.1 we see the transition rates only depend on the population densities $q_{(l,s)}$ and so we have a density dependent population process meaning we can construct an ODE for the expected change in the system, which is exact as $n \to \infty$. It is not a given that such a limit exists, however we see that despite the heterogeneity, the fixed proportions and state space (B.1) are well behaved as $n \to \infty$. Briefly, this is because the proportions $q_{(l,s)}$ belong to $[0, 1]$ and the number of groups m and proportions of servers α_l are independent of n.

We now construct the ODE. Since $q_{(l,s)}$ looks at the proportion of servers with queue length s we only need to consider departures from servers with queue length s and arrivals to servers with queue length $s - 1$. Transition rates (B.2) and (B.3) causes an expected change of $-1/n$ and $1/n$ respectively which gives

$$\frac{dq_{(l,s)}}{dt} = \lambda P_{(l,s-1)} - \mu^l \left(q_{(l,s)} - q_{(l,s+1)}\right). \tag{B.7}$$

□