



HAL
open science

Refinements of Mean Field Approximation

Nicolas Gast

► **To cite this version:**

Nicolas Gast. Refinements of Mean Field Approximation. Performance [cs.PF]. Université Grenoble Alpes, 2020. ⟨tel-02509756⟩

HAL Id: tel-02509756

<https://theses.hal.science/tel-02509756v1>

Submitted on 17 Mar 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Habilitation Thesis (HDR)
(Habilitation à Diriger des Recherches)

REFINEMENTS OF MEAN FIELD APPROXIMATION

NICOLAS GAST

Soutenu à Grenoble, le 30 janvier 2020.

RAPPORTEURS

SEM BORST	TU/e Eindhoven and IBM Research
KAVITA RAMANAN	Brown University
PHILIPPE ROBERT	Inria, Paris

EXAMINATEURS

URTZI AYESTA	CNRS, Toulouse
MICHEL BENAÏM	Univ Neuchâtel
NICOLAS CHAMPAGNAT	Inria, Nancy
DENIS TRYSTRAM	Grenoble-INP and Univ. Grenoble Alpes

Abstract – The design of efficient algorithms is closely linked to the evaluation of their performances. My work focuses on the use of stochastic models for the performance evaluation of large distributed systems. I am interested in developing tools that can characterize the emergent behavior of such systems and improve their performance. This leads me to solve stochastic control and optimization problems, notably through operations research methods. These problems suffer from combinatorial explosion: the complexity of a problem grows exponentially with the number of objects that compose it. It is therefore necessary to design models but also algorithmic processes whose complexity does not increase too rapidly with the size of the system. This document summarizes a few of my contributions on the use and the refinement of mean field approximation to study the performance of distributed systems and algorithms.

Résumé La conception d'algorithmes efficaces est étroitement liés à l'évaluation de leurs performances. Mon travail se concentre sur l'utilisation de modèles stochastiques pour l'évaluation des performances des grands systèmes distribués. Je développe des outils permettant de caractériser le comportement émergent de ces systèmes et d'améliorer leur performance. Cela me conduit à résoudre des problèmes d'optimisation et de contrôle stochastique. Ces problèmes généralement souffrent d'un problème d'explosion combinatoire: la complexité d'un problème croît de façon exponentielle avec le nombre des objets qui le composent. Il faut donc concevoir des modèles mais aussi des processus algorithmiques dont la complexité n'augmente pas ou peu avec la taille du système. Ce document résume quelques unes de mes contributions sur l'utilisation et le raffinement des méthodes de type *champ moyen* pour étudier la performances d'algorithmes de contrôle pour des systèmes distribués.

ACKNOWLEDGMENTS

This document is a summary of a part of my research. The path leading to this document has not always been a straight line and I would like to thank here all the people that made this possible.

First, I would like to thank Sem Borst, Kavita Ramanan and Philippe Robert to have accepted to serve as referees (*rapporteurs*) of my habilitation. They dedicated a large amount of their time to provide insightful comments about my work. I would also like to thank Urtzi Ayesta, Michel Benaïm, Nicolas Champagnat and Denis Trystram to have accepted to be part of the jury as *examimateurs*. It is an honor to have been accepted as peer by such renowned scientists.

I would also like to thank three persons that guided me through my early scientific career: Francois Baccelli who I met during my master and who was the first person to introduce me to performance modeling and stochastic processes, and without whom I might have chosen a very different path; Bruno Gaujal, who accepted me as his PhD student and of whom I always admire the enthusiasm and curiosity about almost anything; and Jean-Yves Le Boudec, with whom I spent four rich years at EPFL and who tried to teach me scientific rigor and how to apply theoretical concepts to very practical problems.

Coming to work everyday is a real pleasure thanks to the many colleagues in Grenoble that for most are also close friends (Annie, Arnaud, Bary, Bruno, Denis, Greg, Olivier, Patrick, Pierre, Jean-Marc,...) and thanks to all PhD students or post-docs with whom I have or had the pleasure to work with (Guillaume, Josu, Benoît, Taki, Mouhcine, Vitalii, Henry, Kimang, Chen, Thomas,...). I am also grateful to all my co-authors for fruitful scientific discussions.



Last but not least, all this work would not have been possible without the support of my friends and family. A particular thought goes to my wife Mathilde and our children Elliot, Martin and Lise. Thanks you for being wonderful human beings and for all the joy that you bring, for giving me slack when I want but bringing me back to reality when needed. Thank you for being there.

CONTENTS

1	INTRODUCTION	1
1.1	Context and Motivation	1
1.2	Organization of the Manuscript	2
2	MEAN FIELD METHODS	5
2.1	Construction of the Approximation : an Example	5
2.2	Density Dependent Population Processes	8
2.3	Proof of Accuracy via Stochastic Approximation	10
2.4	Discontinuous Drifts and Differential Inclusions	13
2.5	Extension to Synchronous Population Models	13
3	SMALL SYSTEMS AND REFINED MEAN FIELD	15
3.1	Accuracy of the Classical Mean Field Approximation	15
3.2	System-Size Expansions	20
3.3	Relation with Moment-Closure Approximation	24
3.4	Application to Load Balancing	25
3.5	Extensions and Open Question	31
4	HETEROGENEOUS MODELS AND CACHE ALGORITHMS	33
4.1	The RAND(m) Policies and Variants	33
4.2	Exact Steady-State Analysis	34
4.3	Transient Analysis via Mean Field	36
4.4	A Heterogeneous Refined Approximation	39
4.5	Extensions and Open Questions	42
5	SPATIAL INTERACTIONS AND BIKE-SHARING SYSTEMS	45
5.1	The Bike Sharing Model	45
5.2	Spatial Refinement and Pair Approximation	49
5.3	Forecasting via Mean Field Approximation	52
5.4	Extension and Open Questions	55
6	CONCLUSION AND PERSPECTIVES	57
	References	59

INTRODUCTION

1.1 CONTEXT AND MOTIVATION

The objective of my research is to develop new approximation techniques in order to solve optimization problems that arise in large systems such as computer networks. My contributions ranges from theoretical work on stochastic models and optimization to applications in communication networks or electricity networks. This document summarizes only a part of my research. It is focused on mean field approximation, its understanding and its application.

Good system design needs good performance evaluation, which in turns needs good modelization tools. Many distributed systems can be effectively described by stochastic population models. These systems are composed of a set of objects, agents, or entities interacting with each other. Each individual agent is typically described in a simple way, as a finite state machine with few states. An agent changes state spontaneously or by interacting with other agents in the system. All transitions happen probabilistically and take a random time to be completed. By choosing exponentially distributed times, the resulting stochastic process is a continuous-time Markov chain with a finite state space. The construction of a Markov chain is therefore quite simple. Markov chains, however, suffer from the state-space explosion: the state-space needed to represent a system with N agents grows exponentially with the number of agents. This results in the need for approximation techniques.

One general idea is to try to characterize the emerging behavior of a system by studying the effect that the local interactions between agents have on the global dynamics – similarly to what is done in statistical physics to characterize behavior of interacting particles. To do so, mean field approximation is a very effective technique that can be used to study transient probability distribution or steady-state properties of such systems when the number of entities N is very large. The idea of mean field approximation is to replace a complex stochastic system, that we denote by $\mathbf{X}^{(N)}$, by a simpler deterministic dynamical system, that we denote by \mathbf{x} . This dynamical system is constructed by assuming that the objects are asymptotically independent. Each object is viewed as interacting with an average of the other objects (the mean field). When each object has a finite or countable state-space, this dynamical system is usually a non-linear ordinary differential equation (ODE).

Mean field approximation has become a common tool to study systems of interacting agents. It is used in many domains such as game theory [60], performance analysis and design of distributed systems [67], multi-agent reinforcement learning [93] or neural networks [4]. The rationale behind mean field approximation is that when the number of agents N is large, each individual has a minor influence on the mass. Mean field models are specified by a simple set of ordinary differential equations (or partial differential equation). This approximation is known to be exact as N goes to infinity for certain models, see *e.g.*, [57], [62], [8] but it can be very poor for other models ($N = 10$ to $N = 100$). The accuracy of the estimates provided by mean field approximation varies greatly depending on the systems and a full understanding of when it can be applied remains to be done. This document summarizes some of the research that I did in this context and that can be summarized in two main objectives:

1. Understand when mean field approximation can be applied and what is its accuracy;
2. Use this understanding to propose more accurate approximation methods.

In particular in this document, we will study refinements of the classical mean field approximation that work for: systems with a low to moderate number of objects (Chapter 3), heterogeneous systems (Chapter 4), and systems with spatial interactions (Chapter 5). Each time, the proposed approximations are validated by using interesting case studies.

1.2 ORGANIZATION OF THE MANUSCRIPT

This manuscript synthesizes my research contributions to the application, development and refinement of new approximation techniques in the context of performance evaluation of large systems. The manuscript is organized in several chapters that reflects the objectives mentioned above.

1. Chapter 2 covers the basic material on **mean field approximation**. Starting with an example, we review a few classical models that are used in the literature. We will present the basic convergence results associated with mean field approximation by making the link with stochastic approximation algorithms. While this chapter mostly covers related work, we also present an extension that concerns **discontinuous drift** and that is based on [P1], [P11], [P12], [P2].
2. Chapter 3 contains the main technical contribution of the manuscript. We first study the **accuracy of the classical mean field approximation** for a general class of population models, called the

density dependent population processes. We then use this characterization to present two **refined mean field approximation** methods. These methods refine the classical approximation by adding a $O(1/N)$ and a $O(1/N^2)$ expansion terms. We prove that these methods are much more accurate than the classical mean field approximation when the population size N is small to moderate. This chapter is based on [P16], [P23], [P22], [P20], [P24], [P17].

3. Chapter 4 discusses a model of **cache replacement algorithms**. These policies aim at adapting the content of a cache to the popularities of requested objects seen by the cache. This model is an example of a **heterogeneous mean field** models. We use it to illustrate how it is possible to construct a (refined) mean field approximation whose accuracy can be proven. This chapter collects contributions from [P28], [P27], [P15], [P25], [P26].
4. Chapter 5 presents a model of **bike-sharing systems**. We describe such a system by a closed queuing network for which we use mean field approximation to derive key performance indicator and to study the performance of natural re-balancing heuristics. In this chapter, we also illustrate how mean field can be modified to model **spatial interactions**. This spatially refined approximation uses pair approximation. This chapter is based on [P7], [P14], [P6], [P19].

Finally, Chapter 6 concludes the manuscript by presenting future research directions, in particular in the context of optimization and game theory.

This chapter reviews some basic material on mean field approximation. We start in Section 2.1 with an example of cache replacement policies that serves as an illustration of what is mean field approximation and how it can be constructed. This example will be further studied in Chapter 4. This example is an element of a broader class of model called density dependent population process whose definition is recalled in Section 2.2. We then discuss in Section 2.3 how stochastic approximation can be used to study the convergence of such models to their mean field approximation. We present an extension of these results to discontinuous drifts in Section 2.4 and to synchronous population processes in Section 2.5.

While most of the chapter is composed of related work, Section 2.4 is based on [P1], [P11], [P12], [P2].

2.1 CONSTRUCTION OF THE APPROXIMATION : AN EXAMPLE

Caches form a key component of many computer networks and systems. A large variety of cache replacement algorithms has been introduced and analyzed over the last few decades. In this section, we illustrate the construction of a mean field approximation for one of the simplest cache replacement policy called RANDOM [35], [3], [28].

The main purpose of this example is to illustrate two ways to construct mean field approximation: by assuming that objects are independent or by looking a population process. Later in Chapter 4, we will discuss more complicated cache replacement policies, based essentially on our papers [P26], [P25]. In particular, we will relate this approximation with the TTL approximation that was originally introduced in [28] and that has become popular since it has been rediscovered independently in [23].

2.1.1 The RANDOM policy

The RANDOM policies works has follows. We consider a single cache that can contains up to m items. When an item, say k , is requested, two events can occur:

- If the item is not in the cache, it is inserted in the cache by replacing a *randomly* chosen item. We call this a miss.
- If the item is in the cache, it is served directly from the cache and the state of the cache is unmodified. We call this a hit.

The IRM model is often described as a discrete-time model where item k is requested with probability p_k at each time step (independently of the previous requests). The two models are essentially equivalent by setting $p_k = \lambda_k / (\sum_\ell \lambda_\ell)$.

To construct our model, we assume the well-known Independent Reference Model (IRM): there is a collection of N items that can be requested. All items have the same size and the cache can store up to m items. An item k has a popularity λ_k and the requests for item k arrive at the cache according to a Poisson process of intensity λ_k . The request processes are assumed to be independent (hence the name IRM).

The state of the cache can naturally be represented by a Markov process by tracking the content of the cache. By using reversibility, one can show that the stationary measure has a product form [53], [35] and performance metric quantities like miss probabilities can be computed by using dynamic programming techniques [29].

In what follows, we show how to construct a mean field approximation of this policy. When applied to steady-state analysis, this mean field approximation allows us to retrieve the classical approximation presented in [29]. More interestingly, it allows a fast and easy characterization of the transient dynamics that can be easily extended to more complex replacement policies, see Chapter 4.

2.1.2 Construction via Object Independence


Let us write $X_k(t)$ a binary variable that equals 1 if Item k is in the cache and 0 otherwise. An item, say k , that is not in the cache is requested at rate λ_k . In which case it will replace a randomly chosen item that is in the cache. Hence, Item ℓ will be the randomly chosen item with probability $\ell X_\ell(t) / m$. This shows that the process $\mathbf{X} = (X_1, \dots, X_N)$ is a Markov chain whose transitions are, for all $k, \ell \in \{1 \dots N\}$:

$$\mathbf{X} \mapsto \mathbf{X} + \mathbf{e}_k - \mathbf{e}_\ell \text{ at rate } \lambda_k(1 - X_k) \frac{X_\ell}{m},$$

where \mathbf{e}_k is a vector whose k th component equals 1, all other being equal to 0.

Let us now zoom on Item k . The variable X_k evolves as in Figure 2.1(a): if the item is not in the cache ($X_k = 0$), it is requested at rate λ_k and enters the cache. This explains the transition from $X_k = 0$ to $X_k = 1$ at rate λ_k . If Item k is already in the cache, it is ejected when another item that is not in the cache is requested (this occurs at rate $\sum_\ell \lambda_\ell(1 - X_\ell)$) and Item k is the one chosen to be replaced (this occurs with probability $1/m$).

It should be clear that X_k is not a Markov chain in general because the evolution of X_k depends on the state of the other objects X_ℓ . If we make the approximation that all objects are independent, we obtain a new stochastic process $\bar{\mathbf{X}}$ that is described in Figure 2.1(b). Here, the evolution of each object is a Markov chain: the states of two objects are independent and objects are only coupled through their mean behavior. Such an approximation is what is a *mean field approximation*.



(a) Original model for RANDOM (b) Mean field approximation

Figure 2.1: RANDOM policy: evolution of the state of Item k .

The quantities $x_k = \mathbb{E}[\bar{X}_k]$ satisfy the following system of differential equation:

$$\dot{x}_k = -\lambda_k(1 - x_k) + \frac{1}{m} \sum_{\ell} \lambda_{\ell}(1 - x_{\ell})x_k. \quad (2.1)$$

Note that Equation (2.1) has a unique fixed point π that satisfies the following equation:

$$\pi_k = \frac{\lambda_k}{\lambda_k + z}, \text{ where } z \text{ is such that } \sum_{\ell} \pi_{\ell} = m. \quad (2.2)$$

The solution of the above equation is unique and corresponds to the approximation proposed in [26].

For now, we did not justify mathematically why Equation (2.1) should be a good approximation of the true value of $\mathbb{E}[X_k]$. This will be the subject of the following sections. Before that, let us just illustrate how accurate this approximation is. For that, we consider a system with $N = 10$ items that have a Zipf popularity with exponent $1/2$: Item k has a popularity $\lambda_k \propto k^{1/2}$ for $k \in \{1 \dots 10\}$. The cache size is $m = N/2 = 5$. In Figure 2.2, we compare the expected hit probability $\sum_k p_k \mathbb{E}[X_k(t)]$ of the original model (obtained by simulation) with the approximation $\sum_k p_k \mathbb{E}[\bar{X}_k(t)]$. We observe a very good match between the two quantities. Note that the value for $\sum_k p_k \mathbb{E}[\bar{X}_k(t)]$ can be easily computed by numerically integrating the ODE (2.1), the value for $\sum_k p_k \mathbb{E}[X_k(t)]$ has been computed by averaging over 10^4 independent simulations. The fixed point of the mean field approximation, given by (2.2), is also a very good approximation of the steady-state hit probability: the exact value is 0.536 while the fixed point gives 0.533.

It should be clear that the method described above can be extended to interacting objects that have more than two states. This is for example what we did in [P26], [P27], [P25] to analyze variants of the RANDOM policy. This will be further discussed in Chapter 4.

2.1.3 Population Processes

We now describe a mathematical framework that can be used, and for example has been used in [45], [83], [48] to study RANDOM or

More generally, Equation (2.1) is an equation of the form $\dot{\mathbf{x}} = \mathbf{x}Q(\mathbf{x})$, where $Q(x)$ is the generator of a Markov chain. Such an equation describes the evolution of an object that interacts with its probability distribution. This equation is similar to a McKean-Vlasov equation that we found in system of interacting particles [65] and that is now widely used in mean field games [47].




Figure 2.2: Cache: Evolution of the expected hit probability $\sum_k p_k \mathbb{E}[X_k(t)]$ as a function of time. We compare with the mean field approximation $\sum_k p_k \mathbb{E}[\bar{X}_k(t)]$. The system has $N = 10$ objects.

LRU caches. Instead of having N different items, we now assume to have K classes of items and N/K items per class, where all items of Class k have the same popularity λ_k . Let us now write $X_k(t) \in [0, 1]$ the fraction of class k items that are in the cache and assume that the cache can store up to mN/K items. The transitions of this model can be written as:

$$\mathbf{X} \mapsto \mathbf{X} + \frac{K}{N}(\mathbf{e}_k - \mathbf{e}_\ell) \text{ at rate } \frac{N}{K} \lambda_k (1 - X_k) \frac{X_\ell}{m},$$

Note that when $N = K$, this model coincides with the model described in the previous section. In the above equation, increasing N reduces the impact of each transition and accelerates the rate of the transitions. As we show in the next section, this is an example of what is called a density dependent population process.

2.2 DENSITY DEPENDENT POPULATION PROCESSES

Typically, we will use $\mathcal{E} = \mathbb{R}^d$ or $\mathcal{E} = \mathbb{R}^N$ (the set of infinite sequences) with an appropriate norm.

Density-dependent population processes are a popular model of population processes that have been introduced in [56], [57]. We recall the definition here. Such a population process is a sequence of continuous-time Markov chains $(\mathbf{X}^{(N)})$. For each N , $\mathbf{X}^{(N)}$ evolves on a subset of some Banach space $(\mathcal{E}, \|\cdot\|)$. We assume that there exists a set of vectors $\mathcal{L} \in \mathcal{E}$ and a set of functions $\beta_\ell : \mathcal{E} \rightarrow \mathbb{R}^+$ such that $\mathbf{X}^{(N)}$ jumps from \mathbf{X} to $\mathbf{X} + \ell/N$ at rate $N\beta_\ell(\mathbf{X})$ for each $\ell \in \mathcal{L}$. Our caching model of §2.1.3 is a density dependent population processes. The next chapters will contain many other examples.

We will refer to the chain $\mathbf{X}^{(N)}$ as the system of size N , although N does not *a priori* correspond to the system's size. For such a system, we

define the drift f as the function that associates to $\mathbf{x} \in \mathcal{E}$ the quantity $f(\mathbf{x})$ that is defined by

$$f(\mathbf{x}) = \sum_{\ell \in \mathcal{L}} \ell \beta_{\ell}(\mathbf{x})$$

The drift does not depend on N . This definition means that $f(\mathbf{x})dt$ is the expected variation of $\mathbf{X}^{(N)}$ that would start in state \mathbf{x} during a small time interval dt :

$$\mathbb{E}[\mathbf{X}(t+dt) - \mathbf{X}(t) \mid \mathbf{X}(t) = \mathbf{x}] = f(\mathbf{x})dt + o(dt).$$

The mean field approximation of such a population process corresponds to the solution of the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$. In our caching example, this would correspond to Equation (2.1).

This model or close variants has been widely used to study the performance of computer-based systems such as dissemination algorithms [22], SSDs [84], and most importantly the analysis of resource allocation strategies in server farms or data centers: such a system is composed of a large number of servers that interact because of scheduling or routing strategies [P10], [63], [67], [82], [90], [66].

It can be shown that under mild conditions (essentially if the drift f is Lipschitz-continuous) then, the stochastic process $\mathbf{X}^{(N)}$ converges almost surely to \mathbf{x} as N goes to infinity [P2, Theorem 1]. More precisely, if \mathbf{x} denotes the solution starting in $\mathbf{x}(0)$ and if the initial conditions $\mathbf{X}^{(N)}(0)$ converges to $\mathbf{x}(0)$ as N (in probability), then for all T :

$$\lim_{N \rightarrow \infty} \sup_{t \leq T} \left\| \mathbf{X}^{(N)}(t) - \mathbf{x}(t) \right\| = 0 \quad (\text{in probability}). \quad (2.3)$$

When one wants to compute values for a finite time-horizon T , the necessary conditions to apply this result can be done by a syntactic analysis of the model. This is more complicated for steady-state analysis [24], in which case special Lyapunov functions have to be found (see for example [P15]).

This is illustrated in Figure 2.3 in which we consider the same caching model as in Section 2.1.3 with different scaling parameters $N \in \{10, 100, 1000\}$ (recall that this corresponds to N items). We observe that, as predicted by the theory, the trajectories of the stochastic system converge to the mean field approximation.

Note that the data reported in Figure 2.2 corresponds to the case $N = 10$ of Figure 2.3. At first sight, it might seem strange that the agreement between the simulation and the mean field approximation seems better on the former than on the later. This discrepancy is explained by the fact that Figure 2.2 depicts the *expected* hit probability $\sum_k \lambda_k \mathbb{E}[X_k]$ whereas Figure 2.3 depicts the a stochastic trajectory of $\sum_k \lambda_k X_k$. We will comment more on this in the next chapter.




Figure 2.3: Scaled caching model: hit probability $\sum_k \lambda_k X_k$ as a function of time compared to the mean field approximation $\sum_k p_k \mathbb{E}[\bar{X}_k]$. We plot the trajectory of one simulation for systems with N objects and $N \in \{10, 100, 1000\}$.

2.3 PROOF OF ACCURACY VIA STOCHASTIC APPROXIMATION

There are essentially two classes of methods to prove that mean field approximation is accurate. The most classical way is to reason in terms of trajectory and to show that, when some scaling parameter tends to infinity, the behavior of the stochastic converges (almost surely or in probability) to the mean field approximation. This often relies on martingale concentration argument and lead to results like Equation (2.3). It can be used for transient [8], [56] or stationary regime [13]. This can be done by using stochastic approximation. We recall the main principle in Section 2.3.1 and its extension to steady-state behavior in Section 2.3.2.

The bound obtained by this trajectorial analysis are often quite loose compared to what is observed in practice. A second kind of methods is to use ideas inspired by Stein's method and [54], [80], [95] to relate the accuracy of the mean field approximation and the gap between the generators of the Markov chain and of the mean field ODE. This method will be the main tool that we will use in Chapter 3.

2.3.1 Stochastic Approximation

Let us consider a density dependent population process \mathbf{X} . The continuous time Markov chain \mathbf{X} can be viewed as a discrete time Markov chain that we denote \mathbf{Y} such that $\mathbf{X}(t) = \mathbf{Y}_{P(t)}$, where $P(t)$ is a Poisson process of intensity $N\tau$, with $\tau = \sup_{\mathbf{x} \in \mathcal{E}} \sum_{\ell} \beta_{\ell}(\mathbf{x})$. For a vector $\ell \in \mathcal{L}$, the discrete time Markov chain has the following transition:

$$\mathbf{P} \left[\mathbf{Y}_{n+1} = \mathbf{Y}_n + \frac{1}{N} \ell \mid \mathbf{Y}_n \right] = \frac{1}{\tau} \beta_{\ell}(\mathbf{Y}_n).$$

Let us define a sequence of random variables \mathbf{M} by $\mathbf{M}_{n+1} = \mathbf{Y}_{n+1} - \mathbf{Y}_n - \frac{1}{\tau N} f(\mathbf{Y}_n)$. By definition of the drift, \mathbf{M} is a martingale difference sequence, *i.e.*, $\mathbb{E}[\mathbf{M}_{n+1} | \mathcal{F}_n] = 0$. This allows us to write the process \mathbf{Y}_n as a stochastic approximation algorithm with constant step size:

$$\mathbf{Y}_{n+1} = \mathbf{Y}_n + \frac{1}{\tau N} (f(\mathbf{Y}_n) + \mathbf{M}_{n+1}). \quad (2.4)$$

Since its introduction in the 50s, the theory of stochastic approximation algorithms has been thoroughly developed, with a particular emphasis on the relation between the stochastic process defined by the recurrence equation (2.4) and the solution of the dynamical system defined by the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ [5], [12], [58]. To make the connection more explicit, let us denote $\epsilon_{N\tau t} = \frac{1}{N\tau} \sum_{k=1}^{N\tau t} M_{k+1}$. This allows us to write the value of \mathbf{Y} at time $N\tau t$ as

$$\mathbf{Y}_{N\tau t} = \mathbf{Y}_0 + \frac{1}{N\tau} \sum_{k=1}^{N\tau t-1} f(\mathbf{Y}_k) + \epsilon_{N\tau t}. \quad (2.5)$$

Under quite general conditions on \mathbf{M} , it can be shown for any time T , $\sup_{t < T} \|\epsilon_{N\tau t}\|$ converges almost surely to 0 as N goes to infinity. This shows that $\epsilon_{N\tau t}$ is a “small” noise.

A great advantage of studying the almost sure convergence of $\mathbf{Y}_{N\tau t}$ by using the stochastic approximation method Equation (2.4) is that it requires very few assumptions. Equation (2.5) is essentially a Euler discretization of the ordinary differential equation $\dot{\mathbf{x}} = f(\mathbf{x})$ with a noise $\epsilon_{N\tau t}$ that vanishes as N becomes large. Proving that $\mathbf{Y}_{N\tau t}$ is close to $\mathbf{x}(t)$ can therefore be done by studying the Euler discretization and proving that it is numerically stable with respect to small error. This is immediate for example when f is a Lipschitz-continuous function, see for example the classical paper of [8]. This method can also be easily generalized when the function f is not continuous. This is for instance what we have done in [P11] when we showed that when f is not continuous, it suffices to replace the ODE by a differential inclusion $\dot{\mathbf{x}} \in F(\mathbf{x})$, see also Section 2.4. Another important generalization of the method is when the system exhibits multiple time-scales, for instance, if the population process $\mathbf{X}(t)$ evolves in an environment $Z(t)$ that changes state at a much faster rate. In this case, one can again construct an ODE by using a time averaging over the possible values of Z [6].

2.3.2 Steady-State Dynamics

The previous results establish that, under mild condition, the trajectories of a system of interacting objects converges to a deterministic dynamical system $\dot{\mathbf{x}} = f(\mathbf{x})$ over any finite time interval. This result can be extended in the stationary regime but this extension has to be handled with care. More precisely, let us consider that, for each

A sufficient condition for example is if the condition second moment of M_n is bounded: $\text{var}[M_{n+1} | \mathcal{F}_n] \leq \sigma^2$. In this case the convergence is a direct consequence of Kolmogorov's inequality for martingales.

*A point π is called an attractor of the ODE if all trajectories of the ODE converge to π . When f is continuous, such a point is of course a fixed point, *i.e.*, such that $f(\pi) = 0$.*

The Birkhoff center of a dynamical system is defined as the closure of the set of recurrent points. A point x is said to be recurrent if $\liminf_{t \rightarrow \infty} \|\Phi_t(x) - x\| = 0$, where $\Phi_t(x)$ denotes the value at time t of the solution of the ODE that starts in x at time 0.

system size N , the Markov chain \mathbf{X}^N has a unique stationary measure π^N . A classical convergence result on the steady-state convergence of mean field models (see for instance [8], [7]) shows that, if the ODE has a unique attractor π , then the sequence of stationary measures π^N concentrate on π as N goes to infinity.

However, characterizing the asymptotic behavior of a differential equation is often a difficult task. In particular, showing that the differential equation has a unique fixed point does not guarantee this fixed point to be an attractor: the situation is more complicated as there might be limit cycles or chaotic behaviors. The only generic results that can be shown in this case is that the sequence of stationary π^N measure concentrate on the Birkhoff center of the ODE. We refer the reader to [8], [24] for a more thorough discussion of the validity of the fixed point method and its application to a model of the 802.11 MAC protocol.

When one tries to study the performance of a system, the case with a unique attractor is often more convenient because if the ODE has multiple attractors, it is often quite difficult to estimate what is the proportion of the time that the system will spend near one or the other point. A large proportion of the mean field models studied in the performance evaluation community seem to have a unique fixed point that is also an attractor. Yet, while the uniqueness of the fixed point is often proven, the proof of its attractiveness is often missing and left as future work or as a conjecture. While there is no general theory to show that a fixed point is an attractor, there exists a few classical methods that may or may apply depending on the particular model:

- One of the most commonly used method is to find a Lyapunov function. There is no generic method to find a Lyapunov function. Yet, for systems of interacting Markov chains, like the ones we study, it is often possible to use relative entropy to build a Lyapunov function. A (relatively) generic approach is presented in [27], [19]. We have used a variant of this approach in [P7], [P15].
- Many systems can be proven to be monotone. Monotonicity can be sometimes be used to prove the convergence to a fixed point. This has been used for example in the context of load balancing systems in [67], [82], [85].
- A interesting last case is when the original stochastic system is reversible for all system sizes N . It is shown in [14] that under this condition, showing the uniqueness of the fixed point is a sufficient condition. This has been for instance used in [83].

2.4 DISCONTINUOUS DRIFTS AND DIFFERENTIAL INCLUSIONS

Most of the result on stochastic approximation assume that the drift f is Lipschitz-continuous in which case the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ is well defined: it has a unique solution and the trajectories of $\mathbf{X}^{(N)}$ converge to this unique solution. The condition of Lipschitz-continuity limits the applicability of these results in many practical cases, in particular, for systems exhibiting threshold dynamics or with boundary conditions.

In [P11], we study the limiting behavior of population processes that have a discontinuous drift f . As a simple example, consider a simple queuing system with one buffer and a processors that serve packet at rate $2N$. If packets arrive at rate N , and if $\mathbf{X}^{(N)}$ denotes the number of packets in the queue rescaled by N^{-1} , then one obtain a density dependent population process whose drift is $f(x) = -1$ if $x > 0$ and $f(0) = +1$. f is not continuous and the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ has no solution.

A proper way to define solutions of an ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ with non-continuous right-hand side f is to use *differential inclusions* (DI) instead. The ODE is replaced by the following equation

$$\dot{\mathbf{x}} \in F(\mathbf{x}), \quad (2.6)$$

where F is a set-valued mapping, defined as the convex hull of the accumulation points of the drift. In the above example, one would have $F(x) = \{-1\}$ if $x > 0$ and $F(0) = [-1, 1]$.

A differential inclusion may have multiple solutions. The main result of [P11] is that over any finite time interval, the trajectories of $\mathbf{X}^{(N)}$ converges to the set of the solutions of the differential inclusion. We also provide a rate of convergence when the set of solution is one-sided Lipschitz-continuous.

A function is one-sided Lipschitz-continuous with respect to a scalar product $\langle \cdot, \cdot \rangle$ if $\langle f(x) - f(y), x - y \rangle \leq L\|x - y\|^2$. A function is Lipschitz-continuous if in addition $\langle f(x) - f(y), x - y \rangle \geq -L\|x - y\|^2$.

2.5 EXTENSION TO SYNCHRONOUS POPULATION MODELS

So far, we have only presented what we call *asynchronous* mean field models. By asynchronous, we mean that each transition that occurs in the systems only affects a limited number of objects, *i.e.*, objects evolves asynchronously. In this case the mean field approximation is given by a continuous time dynamical system (often a system of ordinary differential equations). This document focuses on this model as it is the most studied case *e.g.* [P2], [8], [56]. A second type of model arises when objects are synchronous. In this case the mean field approximation is a discrete time dynamical system [62], [P8], [81]. In this part, we briefly describe this model, based on a simplified version of the model [62].

The synchronous mean field model of [62] models composed N interacting objects in which each object evolves in a finite state space. Time is *discrete* and $\mathbf{X}(t)$ denotes the empirical measure at time t , *i.e.*,

$X_i(t)$ fraction of objects in state i at time t . The behaviour of the system \mathbf{X} is characterized by a (time homogeneous) discrete time Markov chain. At each time step t , each object performs a transition, that, conditionally on $\mathbf{X}(t)$, is independent of the other states. Denoting $S_n(t)$ the state of object n at time t , we have:

$$\mathbf{P}[S_n(t+1) = j \mid S_n(t) = i \wedge \mathbf{X}(t)] = K_{ij}(\mathbf{X}(t)).$$

It can be shown that for such a system, as the number of objects N goes to infinity, the stochastic process \mathbf{X} converges almost surely to a deterministic dynamical system \mathbf{x} that satisfies $\mathbf{x}(t+1) = \mathbf{x}(t)\mathbf{K}(\mathbf{x}(t))$.

A great advantage of this model is that the numerical evaluation of the dynamical system $\mathbf{x}(t+1) = \mathbf{x}(t)\mathbf{K}(\mathbf{x}(t))$ can be computed very efficiently. This model has thus been widely applied, for instance to fasten some model checking algorithms [61] or to model opinion dynamics [71]. It can be extended to more coupled dynamics with non-mean field dynamics, see for instance the interesting model of probabilistic automata in [46].

In this chapter, we describe a new approximation, called the *refined mean field approximation* that we developed in a series of recent papers [P20], [P16], [P23], [P22]. Compared to the classical mean field approximation, this refined approximation depends on the system size N . It converges to the classical mean field approximation when N goes to infinity but is much more accurate for small system size (say $N = 10$ to $N = 100$).

In the remainder of the chapter, we first present the methodology that we use to characterize how accurate is the classical mean field approximation in Section 3.1. We then explain how we use this characterization to propose a refined mean field approximation in Section 3.2. We explain its relation to moment closure approximation in Section 3.3. As an illustration, we discuss the application of these methods to study load balancing strategies in Section 3.4. Finally, we present an extension of these methods to synchronous mean field models in Section 3.5.

3.1 ACCURACY OF THE CLASSICAL MEAN FIELD APPROXIMATION

As argued in Section 2.3, mean field approximation is known to be asymptotically exact for many systems. For these systems, the fraction of objects in a given state i at time t , $X_i^{(N)}(t)$, converges to a deterministic quantity $x_i(t)$, as the number of objects N goes to infinity. The rate of convergence of $X_i^{(N)}$ to x_i has been studied by several papers, e.g. [8], [56], [95], that show that the expected distance between the stochastic process $\mathbf{X}^{(N)}$ and \mathbf{x} is of the order of $1/\sqrt{N}$:

$$\mathbb{E} \left[\left\| \mathbf{X}^{(N)} - \mathbf{x} \right\| \right] \approx \frac{C}{\sqrt{N}}. \quad (3.1)$$

This result is like a central-limit-theorem for mean field systems: $\mathbf{X}^{(N)}(t)$ is equal to $\mathbf{x}(t)$ plus $1/\sqrt{N}$ times a Gaussian noise [56]. It was originally proven for finite time-horizon and more recently extended to stationary distributions in [95].

Yet, we believe that Equation (3.1) does not fully explain the accuracy of mean field approximation. As an example, we provide in Table 3.1 results obtained by simulation on how the mean field approximation is accurate for the power of two-choice model (sometimes also called shortest queue-2 or SQ(2)). We report the average queue length in steady-state as a function of the number of servers N for $\rho = 0.9$, denoted by m^N , and its mean field approximation, denoted by m^∞ . We

The two-choice model is composed of N identical servers. Each server maintains a separate queue. When a job arrives, two servers are sampled at random and the job is allocated to a the server that has the shortest queue among those two. See also Section 3.4.

observe two facts: first, the mean field approximation is already very accurate for $N = 100$, and second, the error made by the mean field approximation is approximately $4/N$ which decreases much faster than the $1/\sqrt{N}$ suggested by Equation (3.1).

Number of servers (N)	10	100	1000	$+\infty$
Average queue length (m^N)	2.81	2.39	2.36	2.35
Error ($m^N - m^\infty$)	0.45	0.039	0.004	0

Table 3.1: Average queue length in the two-choice model. The values for a finite number of servers N are obtained by simulation. The value for $N = +\infty$ is the mean field approximation.

This discrepancy comes from the fact that the error term m^N is a function of the expected value of $\mathbf{X}^{(N)}$: $m^N = \sum_{i=1}^{\infty} i \mathbb{E}[X_i^{(N)}]$. This fact is quite general in the performance evaluation community where mean field approximation is often used to estimate expected performance metrics, that are often function of $\mathbb{E}[\mathbf{X}^{(N)}]$. We therefore argue that the metric used Equation (3.1) is not tight enough to measure how accurate mean field approximation will be.

This leads us to focus on a different metric. Instead of studying the distance between $\mathbf{X}^{(N)}(t)$ and $\mathbf{x}(t)$ as in Equation (3.1), we study the distance between the expectation of a function of $\mathbf{X}^{(N)}(t)$ and its mean field approximation. As a norm is a convex function, we expect the latter to be smaller than the former. We show in fact that there is an order of magnitude of difference: under mild conditions (essentially the twice-differentiability of the drift function f), this distance is of order $1/N$:

$$O\left(\frac{1}{N}\right) = \left\| \mathbb{E}[\mathbf{X}^{(N)}] - \mathbf{x} \right\| \ll \mathbb{E} \left[\left\| \mathbf{X}^{(N)} - \mathbf{x} \right\| \right] = O\left(\frac{1}{\sqrt{N}}\right). \quad (3.2)$$

This result holds for the transient regime and can be extended to the stationary regime under the same conditions as [95].

We will describe later Section 3.2 how this “4” can be computed by numerical methods.

This result shows that an expectation estimated via mean field approximation is $1/N$ -accurate. This is what is observed in Table 3.1, where the average queue length m^N is approximately equal to $m^\infty + 4/N$. Equation (3.2) also explains the discrepancy between Figure 2.1 and Figure 2.2: the good agreement between the mean field approximation and the simulation in Figure 2.1 is explained because $\left\| \mathbb{E}[\mathbf{X}^{(N)}] - \mathbf{x} \right\|$ is small. On the other hand, Figure 2.2 shows that $\mathbb{E} \left[\left\| \mathbf{X}^{(N)} - \mathbf{x} \right\| \right]$ is still relatively large for $N = 10$.

Note that the metric Equation (3.2) has also been studied in [9]. In this paper, the authors also obtain a $O(1/N)$ rate of convergence for another kind of mean field models. They call this metric the *bias* of the approximation.

3.1.1 Main Results

We now describe in more details the precise model under which the above result holds. We consider population processes described by the classical model of density-dependent population process that was described in Section 2.2, that is, for each N , $\mathbf{X}^{(N)}$ defines the evolution of the stochastic process of size N and there exists a set of vectors $\mathcal{L} \in \mathcal{E}$ and a set of functions $\beta_\ell : \mathcal{E} \rightarrow \mathbb{R}^+$ such that $\mathbf{X}^{(N)}$ jumps from \mathbf{x} to $\mathbf{x} + \ell/N$ at rate $N\beta_\ell^N(\mathbf{x})$ for each $\ell \in \mathcal{L}$. The drift of the system of size N is f^N , defined by $f^N(\mathbf{x}) = \sum_{\ell \in \mathcal{L}} \ell \beta_\ell^N(\mathbf{x})$. We will denote by $f(\mathbf{x}) = \lim_{N \rightarrow \infty} f^N(\mathbf{x})$ the limit if it exists.

The $1/N$ -accuracy of the mean field approximation holds under essentially two main assumptions:

(A0) $\mathbf{X}^{(N)}$ is a density dependent population process on a Banach¹ space \mathcal{E} such that $\sup_{\mathbf{x} \in \mathcal{E}} \sum_{\ell \in \mathcal{L}} \beta_\ell^N(\mathbf{x}) \|\ell\|^2 < \infty$.

(A1) The drift f is C_2 (i.e., differentiable twice with a continuous second derivative).

Note that (A0) is essentially a technical assumption that ensures that the models that we study are well-defined. Assumption (A1) deserves more discussion. Assuming that the drift is differentiable once is quite classical to obtain mean field convergence (recall Section 2.3) even if some results can be obtained for discontinuous drifts [P11]. Assuming that the drift is *twice* differentiable is a more restrictive condition. It is often verified by interaction models in which independent tuples of agents can interact. It is often not verified when a centralized controller is involved, as for example in the bandit optimal model of [88]. Yet, this condition is almost necessary as we discuss in Section 3.1.3.

Using this result we obtain the following theorem:

Theorem 1 (Based of Theorems 3.1 and 3.3 of [P16]). *Assume that $\mathbf{X}^{(N)}$ satisfies (A0) and (A1) and that there exists a constant $c > 0$ such that $\|\mathbb{E}[\mathbf{X}^{(N)}(0)] - \mathbf{x}\| \leq c/N$ and $\sup_{\mathbf{x} \in \mathcal{E}} \|f^N(\mathbf{x}) - f(\mathbf{x})\| \leq c/N$. Then, for each t , there exists a constant $C(t) < \infty$ such that*

$$\left| \mathbb{E} \left[\mathbf{X}^{(N)}(t) \right] - \Phi_t \mathbf{x} \right| \leq \frac{C(t)}{N}. \quad (3.3)$$

This theorem can be extended to the stationary regime under two additional assumptions:

(A2) For each N , the stochastic process has a unique stationary distribution.

In particular (A1) implies that the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ is well-defined and has a unique solution starting from any initial condition. In what follows we will denote by $\Phi_t \mathbf{x}$ the value of this solution at time t .

¹ In practice, the state spaces \mathcal{E} that we use are mostly subsets of \mathbb{R}^d for $d < \infty$. Yet, it is sometimes convenient to consider infinite dimensional spaces, for instance when considering a mean field limit of a queuing network in which queues are unbounded (see for instance §3.4.1).

(A3) The ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ has a unique attractor π and this attractor is exponentially stable, *i.e.* there exists $a, b > 0$ such that for all $\mathbf{x} \in \mathcal{E}$ and $t > 0$: $\|\Phi_t \mathbf{x} - \pi\| \leq ae^{-bt}$.

Again, Assumption (A2) is a very natural assumption that is needed for our theorem to be stated. Assumption (A3) contains two sub-assumptions: the existence of a unique attractor and its exponential stability. As pointed out in Section 2.3.2, the existence of a unique attractor is necessary to apply the classical concentration results of for instance [8]. The exponential stability is slightly more restrictive but is generally verified in practice. Moreover, as opposed to the global stability, the exponential stability is in general easy to verify numerically (for a finite dimensional space, it only depends on the eigenvalues of the Jacobian at π). This means that our conditions are not more difficult to verify than the ones of [8].

Theorem 2. *In addition to the assumptions of Theorem 1, assume (A2) and (A3). Then the constant $C(t)$ in Theorem 1 is uniformly bounded in time. In particular, there exists a constant $C < \infty$, such that:*

$$\left| \mathbf{E}^{(N)} \left[\mathbf{X}^{(N)} \right] - \pi \right| \leq \frac{C}{N} \quad (3.4)$$

where $\mathbf{E}^{(N)}[\cdot]$ denotes the expectation with respect to the stationary distribution of the system of size N .

3.1.2 The Generator Method

We will not give a full proof of the above theorem but rather present the theory behind and give the main ingredient of the proof. The starting point of our result is to use a generator method similar to the one developed in [54], [95] in a similar mean field context. There are two main ingredients in the proof: The first is to show that the generators of the stochastic system and the ODE are close; The second is to translate this into a bound on the gap between $\mathbb{E}[h(\mathbf{X}^{(N)}(t))]$ and $h(\Phi_t \mathbf{x})$. While the first one is relatively straightforward, the second comes from a careful analysis how the solution of an ODE varies as a function of its initial condition. This is where we use the assumption (A1) of twice-differentiability. For the transient regime, this method is close to the one used in [54]. For stationary regime, this method can be related to the use of Stein's method [80] whose use has been recently popularized in [43], [17], [18]. This methodology has for instance then been used in [95], [96] to establish the rate of convergence of stochastic processes to their mean field approximation, in light or heavy traffic.

Let us denote by $L^{(N)}$ the generator of the system of size N and by Λ the generator of the ODE. To a function h that is differentiable,

these generators associates respectively the functions $L^{(N)}h$ and Λh that are defined as

$$\begin{aligned} (L^{(N)}h)(\mathbf{x}) &= \sum_{\ell \in \mathcal{L}} N\beta_\ell(\mathbf{x})\left(h\left(\mathbf{x} + \frac{\ell}{N}\right) - h(\mathbf{x})\right) \\ (\Lambda h)(\mathbf{x}) &= \sum_{\ell \in \mathcal{L}} \beta_\ell(\mathbf{x})Dh(\mathbf{x}) \cdot \ell = Dh(\mathbf{x}) \cdot f(\mathbf{x}), \end{aligned}$$

where $Dh(\mathbf{x})$ is the derivative of the function h , evaluated in \mathbf{x} .

The first idea of the proof is to remark that if the generator Λ is a first-order Taylor expansion of the generator $L^{(N)}$. More precisely, if a function g is twice-differentiable, then $L^{(N)}g$ converges to Λg at rate N^{-1} . Indeed:

$$\begin{aligned} (L^{(N)} - \Lambda)g\mathbf{x} &= \sum_{\ell \in \mathcal{L}} \beta_\ell(\mathbf{x}) \left(N\left(g\left(\mathbf{x} - \frac{\ell}{N}\right) - g(\mathbf{x})\right) - Dg(\mathbf{x}) \cdot \ell \right) \\ &= \frac{1}{2N} \sum_{\ell \in \mathcal{L}} \beta_\ell(\mathbf{x}) D^2g \cdot (\ell \otimes \ell) + O(1/N^2). \end{aligned} \quad (3.5)$$

The second step of the proof is to use the following equality (referred as a ‘‘standard trick’’ in [54]) to obtain:

$$\mathbb{E} \left[\mathbf{X}^{(N)}(t) - \Phi_{t\mathbf{x}} \right] = -\mathbb{E} \left[\int_0^t (\Lambda - L^{(N)})\Phi_s(\mathbf{X}^{(N)}(t-s)) ds \right]. \quad (3.6)$$

Theorem 1 follows from applying Equations (3.5) with $g = \Phi_t$ to (3.6). To show that, one needs to show that Φ_s is twice-differentiable function which is a classical result in analysis.

The proof of Theorem 2 consists in looking at the limit of Equation 3.6 as t goes to infinity. This equation is replaced by

$$\mathbb{E} \left[\mathbf{X}^{(N)} - \pi \right] = -\mathbb{E} \left[(\Lambda - L^{(N)}) \int_0^\infty \Phi_s(\mathbf{X}^{(N)}) ds \right], \quad (3.7)$$

where the above expectation represents the expectation with respect to the stationary measure of $\mathbf{X}^{(N)}$.

Theorem 2 then follows by using Equations (3.5) with $g = \int_0^\infty \Phi_t dt$. Again, a careful analysis of the dynamical system shows that this function is twice differentiable when this ODE has an exponentially stable attractor, see for example [P16], [95].

3.1.3 Necessity of Twice-Differentiability

The main assumption that is used to obtain the rate of convergence of $1/\sqrt{N}$ of Equation (3.1) is to have a Lipschitz-continuous drift. On the other hand, our results Theorem 1 and Theorem 2 require the drift to be *twice* differentiable. It is quite natural to wonder about the necessity of this assumption.

In [P16], we relax slightly the assumption of twice-differentiability and show that when the derivative of the drift is α -Hölder continuous,

Recall that $\Phi_{t\mathbf{x}}$ is the solution of the equation at time t that starts in \mathbf{x} . The twice-differentiability of Φ_t means that the solution of an ODE is twice-differentiable function of its initial condition.

Equation (3.6) is obtained by studying the function $r(s) = \mathbb{E}[\Phi_s(\mathbf{X}^{(N)}(t-s))]$: the right hand side of Equation (3.6) is $r(0) - r(t)$ while the left hand side is $\int_0^t r'(s) ds$.

The rate of $1/\sqrt{N}$ is also true with a one-sided Lipschitz-continuity, see e.g. [P11].

Recall that a function is α -Hölder continuous if there exists a constant L such that for all \mathbf{x}, \mathbf{y} : $\|f(\mathbf{x}) - f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|^\alpha$. The case $\alpha = 1$ corresponds to Lipschitz-continuity.

then the convergence occurs at rate $O(1/\sqrt{N}^{1+\alpha})$. We also provide a numerical example that shows that in general this exponent is tight and that having a Lipschitz-continuous drift is not sufficient to obtain Equation (3.2). This implies that only Lipschitz-continuity is not sufficient to guarantee a $1/N$ -accuracy of the mean field estimates.

3.2 SYSTEM-SIZE EXPANSIONS

Theorem 1 and 2 provides bounds that guarantees that an expected value is *at most* from a distance of order $1/N$ of its mean field approximation. In [P23], [P20], we show that this rate of convergence is exact. More precisely, we show that under assumptions (A0) and (A1), for any $t > 0$, there exists a constant $V(t)$ such that

$$\lim_{N \rightarrow \infty} N(\mathbb{E}[\mathbf{X}^{(N)}(t)] - \Phi_t \mathbf{x}) = V(t). \quad (3.8)$$

Under assumptions (A2) and (A3) the result also holds for the stationary regime ($t = +\infty$).

While replacing the inequalities of the Theorems 1 and 2 is not difficult, the major contributions of [P23], [P20] is to show that $V(t)$ can be easily computed numerically: when $t > 0$ this constant satisfies a (time-varying) linear differential equation. For $t = +\infty$, this constant is the solution of a linear system. The *refined* approximation that we propose consists in approximating $\mathbb{E}[\mathbf{X}^{(N)}(t)]$ by its mean field approximation plus the $1/N$ term:

$$\mathbb{E}[\mathbf{X}^{(N)}(t)] \approx \Phi_t \mathbf{x} + \frac{1}{N} V(t). \quad (3.9)$$

By using several examples, we show in that this refined approximation is remarkably accurate, even for small system sizes such as $N = 10$. We will provide numerical examples illustrating this in Section 3.4.

Note that Equation (3.9) can be viewed as a first order Taylor expansion of the classical mean field approximation. We also show in [P20] that it is possible to obtain a second order expansion $\mathbb{E}[\mathbf{X}^{(N)}(t)] \approx \Phi_t \mathbf{x} + \frac{1}{N} V(t) + \frac{1}{N^2} A(t)$, and characterize a differential equation satisfied by $A(t)$. To ease the presentation, in the part below, we describe briefly how the first term $V(t)$ can be derived, by using the moment method that we developed in [P20].

3.2.1 Derivation of the Refined Mean Field Approximation

Before stating the main theorem and giving an idea of the proof, let us introduce some important notations. In the proof bellow, the quantities $V(t), W(t) \dots$ denote time-varying tensors but we will drop the dependence on t and simply write V, W, F, Q in order to simplify notation. For i, j, k , the quantities F_j^i and F_{jk}^i denote the first and second

derivative of the i th component of the drift f evaluated at $\Phi_t \mathbf{x}$ (recall that $\Phi_t \mathbf{x}$ is value at time t of the solution of the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$):

$$F_j^i = \frac{\partial f_i}{\partial x_j}(\Phi_t \mathbf{x}) \quad \text{and} \quad F_{jk}^i = \frac{\partial^2 f_i}{(\partial x_j)(\partial x_k)}(\Phi_t \mathbf{x})$$

Similarly, H_j and H_{jk} denote the first and second derivatives of a function h evaluated at $\Phi_t \mathbf{x}$. Last we define the time varying tensor Q as

$$Q^{ij} = \sum_{\ell \in \mathcal{L}} \beta_\ell(\Phi_t \mathbf{x}) \ell_i \ell_j.$$

Using this notation, the constant expressed by Equation (3.8) can be restated in the following theorem, that also shows the equations satisfied by the constants V and W .

Theorem 3 (Theorem 1 of [P20]). *Assume that the model satisfies (A0) and (A1) and that $\mathbf{X}^{(N)}(0) = \mathbf{x}$. Then there exists time-varying tensors $V(t)$ and $W(t)$ such that for any differentiable function h :*

$$\lim_{N \rightarrow \infty} N(\mathbb{E} [h(\mathbf{X}^{(N)}(t))] - h(\Phi_t \mathbf{x})) = \sum_j H_j V^j(t) + \frac{1}{2} \sum_{jk} H_{jk} W^{jk}(t). \quad (3.10)$$

Moreover, V and W satisfy the following set of linear differential equations:

$$\dot{V}^i = \sum_j F_j^i V^j + \frac{1}{2} \sum_{jk} F_{jk}^i W^{jk} \quad (3.11)$$

$$\dot{W}^{ij} = \sum_k (F_k^i W^{kj} + F_k^j W^{ki}) + Q^{ij}. \quad (3.12)$$

If the model also satisfies (A2) and (A3), then the above differential equations (3.11-3.12) have a unique fixed point and Equation (3.10) holds uniformly for all time $t \in \mathbb{R}^+ \cup \{\infty\}$.

Note that in [P20], we also derive a second order expansion in $1/N^2$. The theorem below states the existence of the constants corresponding to the second order expansion. To compute the $1/N$ expansion of Theorem 3, the above differential equation depends on the first and second derivatives of the drift. In the theorem below, we do not present the exact form of the differential equation but it is to be noted that in order to compute the second order expansion in $1/N^2$, one need to consider up to the fourth derivative.

Theorem 4 (Theorem 1 of [P20]). *In addition of the assumptions of Theorem 3, assume that the drift and the function h are four times differentiable. Then there exists time-varying tensors $A(t)$, $B(t)$, $C(t)$ and $D(t)$ of respec-*

tive sizes d , $d \times d$, $d \times d \times d$ and $d \times d \times d \times d$ such that for any time $t > 0$:

$$\begin{aligned} & \lim_{N \rightarrow \infty} N^2 \left(\mathbb{E} \left[h(\mathbf{X}^{(N)}(t)) \right] - \Phi_t \mathbf{x} - \frac{1}{N} \left(\sum_j H_j V^j(t) + \frac{1}{2} \sum_{jk} H_{ij} W^{jk}(t) \right) \right) \\ &= \sum_j H_j A^j(t) + \frac{1}{2} \sum_{jk} H_{jk} B^{jk}(t) + \frac{1}{6} \sum_{jkl} H_{jkl} C^{jkl}(t) + \frac{1}{24} \sum_{jklm} H_{jklm} D^{jklm}(t) \end{aligned}$$

These tensors satisfy time-dependent ordinary differential equation similar to that of V and W of Theorem 3. Moreover, if the model also satisfies (A2) and (A3), then these differential equations have a unique fixed point and the above equation holds uniformly for all time $t \in \mathbb{R}^+ \cup \{\infty\}$.

The proof of Theorem 3 and Theorem 4 that we provide in [P20] is decomposed in two parts. First, we establish the existence of the above constants. This is similar to what we presented in Section 3.1.2 and we will not recall here. In a second part, we show that the constant satisfies the above equation. This second part is more original and we present it here. For the sake of conciseness, the only present the derivation of the first order equation Theorem 3. We refer to [P20] for a full derivation of the second order expansion.

The main idea of the proof of [P20] is to compute the derivatives of the moments of $\mathbf{X}^{(N)} - \Phi_t \mathbf{x}$ with respect to time. For the first moment,

$$\frac{d}{dt} \mathbb{E} \left[X_i^{(N)}(t) - (\Phi_t(\mathbf{x}))_i \right] = \mathbb{E} \left[f_i(\mathbf{X}^{(N)}(t)) - f_i(\mathbf{x}) \right] \quad (3.13)$$

Recall that $\mathbf{X}^{(N)}$ jumps from \mathbf{X} to $\mathbf{X} + \ell/N$ at rate $N\beta_\ell(\mathbf{x})$ that that $f(\mathbf{x}) = \sum_\ell \ell \beta_\ell(\mathbf{x})$.

Applying Equation 3.10 to the above equation shows that

$$\lim_{N \rightarrow \infty} N \mathbb{E} \left[f_i(\mathbf{X}^{(N)}(t)) - f_i(\mathbf{x}) \right] = \sum_j F_j^i V^j + \frac{1}{2} \sum_{jk} F_{jk}^i W^{jk}.$$

This gives Equation (3.11).

For the second moment, let us simplify the notation by denoting $\mathbf{x} = \Phi_t \mathbf{x}$. The evolution of the stochastic process $(X_i^{(N)}(t) - (\Phi_t(\mathbf{x}))_i)(X_j^{(N)}(t) - (\Phi_t(\mathbf{x}))_j)$ can be decomposed in two parts. On the one hand, it jumps from $(X_i - x_i)(X_j - x_j)$ to $(X_i + \ell_i/N - x_i)(X_j + \ell_j/N - x_j)$ at rate $\beta_N(\mathbf{X})$. This translates into a change of

$$\begin{aligned} & (X_i + \frac{\ell_i}{N} - x_i)(X_j + \frac{\ell_j}{N} - x_j) - (X_i - x_i)(X_j - x_j) \\ &= \frac{1}{N}(X_i - x_i)\ell_j + \frac{1}{N}(X_j - x_j)\ell_i + \frac{1}{N^2}\ell_i\ell_j. \end{aligned}$$

Summing over all possible transitions leads to an average variation of

$$(X_i - x_i)f_j(\mathbf{X}) + (X_j - x_j)f_i(\mathbf{X}) + \frac{1}{N}\beta_\ell(\mathbf{X})\ell_i\ell_j$$

On the other hand, $\Phi_t \mathbf{x}$ satisfies a differential equation with drift f , that leads to a variation of

$$-f_i(\mathbf{x})(X_j - x_j) - (X_i - x_i)f_j(\mathbf{x})$$

Summing the above two terms shows that

$$\begin{aligned} & \frac{d}{dt} \mathbb{E} \left[(X_i^{(N)}(t) - (\Phi_t(\mathbf{x}))_i)(X_j^{(N)}(t) - (\Phi_t(\mathbf{x}))_j) \right] \\ &= \mathbb{E} \left[(X_i - x_i)(f_j(\mathbf{X}) - f_j(\mathbf{x})) + (X_j - x_j)(f_i(\mathbf{X}) - f_i(\mathbf{x})) + \frac{1}{N} \beta_\ell(\mathbf{X}) \ell_i \ell_j \right]. \end{aligned}$$

As for Equation (3.11), applying Equation 3.10 leads to the above equation leads after some computation to Equation (3.12).

From this point, it should be clear that an extension of the above method to higher order moments is possible but will lead to much more complicated derivations. In [P20], we derive the expressions for the third and fourth moments which are necessary to obtain the $1/N^2$ term of the expansion. This derivation was made possible by using compact tensor notations and Einstein summation convention (summation over a set of repeated indices). Thanks to this method, we obtain a characterization of the $1/N^2$ constant that, despite being hard to interpret (see the equations for A, B, C, D in [P20, Theorem 1]), can be easily solved numerically.

3.2.2 Numerical Implementation

In order to ease the applicability of the above method, we implemented a numerical library in Python that constructs and solves the above equation. The library is available at https://github.com/ngast/rmf_tool/ [P18]. It takes as an input a description of the model and uses symbolic differentiation to construct the derivatives of the drift and of the functions Q and R .

For the transient analysis, the computation of $V(t), W(t), \dots$ requires a numerical integration of an ordinary differential equation. To do so, we use the function tensor manipulation routines of the library `numpy` (like `tensor.dot` or `reshape`) to transform our complex tensor equation into `numpy` arrays. We then use the function `integrate.solve_ivp` of the library `scipy` [50] to numerically integrate the ODEs for computing $V(t)$ and $W(t)$ of Theorem 1. For the steady-state analysis, the tool uses the python library `scipy.sparse` to construct a sparse system of linear equations and the function `scipy.sparse.linalg.lgmres` to solve the sparse linear system.

A detail analysis of the computation time taken by the method is performed in [P20, §4.3.4 and Figure 1]. In brief, it shows that the computation of the $1/N$ -terms $V(t)$ and $W(t)$ can be done for models with hundreds of dimensions in less than 10 seconds. With the same constraints of 10 seconds, the $1/N^2$ -terms can be computed for models with a few tens of dimensions. Moreover, computing the time-varying constants for the transient regime is more costly than solving the fixed point equations: For a given computation-time budget, one can compute the steady-state constants for a system of doubled size.

3.3 RELATION WITH MOMENT-CLOSURE APPROXIMATION

An alternative way to derive the refined mean field approximation is to reason in terms of moment closure approximation. Moment closure techniques are used in numerous papers in the theoretical biology literature, e.g. [76], [2], [38], [78], [74], [42]. We refer to [55] for a recent review on the subject. In [39], it is argued that the moment closure of order 2 provides the same $O(1/N)$ term as the refined mean field. This relates to a previous result of the same authors [40] where the authors develop an approximation similar to the one of [P23].

Let us consider a density dependent population process \mathbf{X} . As written in Equation (3.13), the derivation of the expectation of $\mathbf{X}^{(N)}$ satisfies:

$$\frac{d}{dt} \mathbb{E} [\mathbf{X}(t)] = \sum_{\ell \in \mathcal{L}} \ell \mathbb{E} [\beta_{\ell}(\mathbf{X}(t))] = \mathbb{E} [f(\mathbf{X}(t))]. \quad (3.14)$$

Note that $f(\mathbf{x})$ is a polynomial of degree 2 in all the examples of the manuscript except for the power of d choices with $d \geq 3$ where f is a polynomial of degree d .

The above equation is exact but if f is not a linear function, the above differential equation is not closed, because it is impossible to express $\mathbb{E}[f(\mathbf{X}(t))]$ as a function of $\mathbb{E}[\mathbf{X}(t)]$. For example if f is a polynomial of degree 2, $f(\mathbf{X}) = \mathbf{A}\mathbf{x} + \mathbf{x}^T \mathbf{B}\mathbf{x}$, the expectation $\mathbb{E}[f(\mathbf{X}(t))]$ involves the second moment of \mathbf{X} .

There are two solutions to close the ODE (3.14). The first leads to the mean field approximation and consists in approximating $\mathbb{E}[f(\mathbf{X}(t))]$ by $f(\mathbb{E}[\mathbf{X}(t)])$. The second solution is to consider higher moments of $\mathbf{X}(t)$. For instance, the second moment evolves as:

$$\begin{aligned} \frac{d}{dt} \mathbb{E} [\mathbf{X}(t)\mathbf{X}(t)^T] &= \sum_{\ell \in \mathcal{L}} \mathbb{E} \left[\left(\mathbf{X}(t)\ell^T + \ell\mathbf{X}(t)^T + \frac{1}{N}\ell\ell^T \right) \beta_{\ell}(\mathbf{X}(t)) \right] \\ &= \mathbb{E} \left[\mathbf{X}(t)f(\mathbf{X}(t))^T + f(\mathbf{X}(t))\mathbf{X}(t)^T + \frac{1}{N}Q(\mathbf{X}(t)) \right], \end{aligned} \quad (3.15)$$

where $Q(\mathbf{x})$ is defined as $Q(\mathbf{x}) := \sum_{\ell} \ell\ell^T \beta_{\ell}(\mathbf{x})$.

Again, the above equations is exact but is not closed if f is not a linear function. For instance, if f is a polynomial of order 2, the right hand side of this equation involves the third moment of $\mathbf{X}(t)$. This process can be continued: the derivative with respect to time of the k th moment will depend on the $(k+1)$ th moment.

The idea of moment closure approximation is to artificially close the equation by approximating a moment of order $k+1$ by lower order moments. A natural way to do this is to assume that the $(k+1)$ th central moment is equal to 0. For instance, the moment closure of order 2 consists in approximating the third central moment by 0: for any i, j, k :

$$\mathbb{E} [(X_i - \mathbb{E}[X_i])(X_j - \mathbb{E}[X_j])(X_k - \mathbb{E}[X_k])] \approx 0.$$

Said otherwise, we obtain a moment closure of order 2 by replacing in Equation (3.15) the terms $\mathbb{E}[X_i X_j X_k]$ by

$$\mathbb{E}[X_i X_j] \mathbb{E}[X_k] + \mathbb{E}[X_i X_k] \mathbb{E}[X_j] + \mathbb{E}[X_j X_k] \mathbb{E}[X_i] - 2\mathbb{E}[X_i] \mathbb{E}[X_j] \mathbb{E}[X_k]$$

The relation between moment closure and our refined mean field approximation is not *a priori* straightforward. In fact, our approach can be used to justify that moment closure provides estimates that are more accurate than the classical mean field approximation. To see that, let us compute what we lose by neglecting the third central moment $\mathbb{E}[(X - \mathbb{E}[X])] \approx 0$. The third central moment is equal to:

$$\begin{aligned} \mathbb{E}[(X - \mathbb{E}[X])^3] &= \mathbb{E}[(X - x)^3] + 3\mathbb{E}[(X - x)^2](x - \mathbb{E}[X]) \\ &\quad + 3\mathbb{E}[(X - x)](x - \mathbb{E}[X])^2 + (x - \mathbb{E}[X])^3 \end{aligned}$$

Denoting by x the mean field approximation, Theorem 3 shows that $\mathbb{E}[X] - x = O(1/N)$, $\mathbb{E}[(X - x)^2] = O(1/N)$ and $\mathbb{E}[(X - x)^3] = O(1/N^2)$. This shows that the third central moment is of order $1/N^2$ which implies that neglecting the third central moment provides an approximation that is $1/N^2$ -accurate (approximation 2MA of [39]).

Similarly, it can be shown that the fourth central moments is of order $O(1/N^2)$ and that the fifth central moment is of order $O(1/N^3)$. Hence, to obtain an approximation that is an order of magnitude better ($1/N^3$ -accurate), one should not neglect the fourth moment but only the fifth moment. Neglecting the fourth moment (approximation 3MA of [39]) leads to a approximation that is $O(1/N^2)$ -accurate, similarly to the approximation 2MA.

3.4 APPLICATION TO LOAD BALANCING

In this section, we show how this refined mean field approximation can be applied to study load balancing strategies. In particular, we show that it is much more accurate than the classical mean field approximation, especially for small system sizes and that it allows to distinguish models that could not be distinguished before, such as the impact of choosing with or without replacement (see Section 3.4.4).

The analysis of load balancing strategies in server farms is an important area of application of mean field approximation. Such a system is composed of a large number of servers that interact because of scheduling or routing strategies [P10], [63], [67], [82], [90], [66]. A typical example is the *power of two-choice*: Mean field approximation has been used in [67], [90] to show that, with an infinite number of servers N , routing a task to the least loaded of two randomly sampled servers significantly reduces the response time compared to a purely random allocation. In this part, we revisit this problem by looking at what happens when the number of servers N is not infinite.

3.4.1 Load Balancing Strategies

We consider a system composed of N identical servers. Jobs arrive at the system according to a Poisson process of rate ρN . The service time of each job is exponentially distributed of mean 1. In this part, we will

compare different strategies to allocate jobs to servers. These policies are often called *load balancing* strategies as their goal is to balance the load among servers. In the remaining of this section, we will compare two resource allocation strategies:

- *d*-choice policy of [67], [90] – For each incoming job, two servers are picked at random. The jobs is allocated to the server with the smallest queue size and are then processed in a “first come first served” order.
- Pull-push model of [66] – Jobs are allocated purely at random to servers. When a server is idle, it asks an other server at random and *pull* one job for this queue if this server has 2 or more jobs. This takes a random time that is exponentially distributed of parameter r .

Our goal is particularly to understand, for a given system size N and load ρ , what is quantitatively the impact of choosing one or the other strategy. The classical mean field approximation does not allow one to distinguish between different system size. We will show that our refined approximation does.

3.4.2 Model and Mean Field Approximation

We first describe in more details the model and its classical mean field approximation. Let $X_i^{(N)}(t)$ denote the fraction of servers with queue size at least i at time t . As the role of each server is symmetrical, $\mathbf{X}^{(N)}(t)$ is a Markov chain whose transitions are as follows. A departure from a server with $i \geq 1$ jobs modifies \mathbf{X} into $\mathbf{X} - N^{-1}\mathbf{e}_i$ and occurs at rate $N(X_i - X_{i+1})$. A pull from a queue with i jobs modifies \mathbf{X} into $\mathbf{X} + N^{-1}(-\mathbf{e}_i + \mathbf{e}_1)$ and occurs at rate $r(1 - X_0)(X_i - X_{i+1})$. An arrival at a queue with i jobs modifies \mathbf{X} into $\mathbf{X} + N^{-1}\mathbf{e}_i$ and occurs at a rate depending on the allocation strategy. If the arrival are routed purely at random, this occurs at rate $N\rho(X_{i+1} - X_i)$. If an arrival is routed to the least loaded of d servers, then it depends whether the servers are picked with or without replacement. With replacement, this occurs at rate $N\rho(X_{i+1}^d - X_i^d)$.

The mean field approximation are the solution of the ODE $\dot{\mathbf{x}} = f(\mathbf{x})$ where the drifts are: for the *d*-choice model:

$$f_i(\mathbf{x}) = \rho(x_{i-1}^d - x_i^d) + (x_{i+1} - x_i). \quad (3.16)$$

For the push-pull models:

$$f_i(\mathbf{x}) = \begin{cases} \rho(x_{i-1} - x_i) + (x_{i+1} - x_i) - r(x_{i-1} - x_i)(1 - x_1) & \text{if } i \geq 2 \\ \rho(1 - x_1) + (x_2 - x_1) + rx_1(1 - x_1) & \text{if } i = 1. \end{cases}$$

Note that for $\rho < 1$ and $r > 0$, these two models have a unique stationary fixed to which all trajectories converge. This can be shown

by monotonicity arguments as in [67], [66] or also by using a Lyapunov function as in [67].

3.4.3 Accuracy of the Refined Mean Field

As said earlier, the mean field approximation is asymptotically exact as the number of servers N goes to infinity. This has been used in the literature to predict what will be the average queue length of such a system, which for these models is $\sum_{i \geq 1} \mathbb{E}[X_i]$. As an illustration, we show in Table 3.2 the expected queue length predicted by the mean field approximation, *i.e.*, the value of $\sum_{i \geq 1} \pi_i$ where π is the fixed point of the corresponding policy.

Load	$\rho = 0.8$	$\rho = 0.9$	$\rho = 0.95$
Expected queue length (push-pull)	1.3333	1.6364	1.8095
Expected queue length (two-choice)	1.5579	2.3527	3.2139

Table 3.2: Expected Queue Length Predicted by the Mean Field Approximation for the two load balancing policies. For the push pull policy, we choose a pooling rate of $r = 1/(1 - \rho)$.

The values presented in Table 3.2 represent the theoretical that one would observe in such a system if the number of servers was infinite. For small system size, these numbers can be quite far from the reality. For instance, a simulation of the push-pull model for $N = 10$ and $\rho = 0.9$ suggests that the expected queue length is closer to 2.30 which is 45% larger than the 1.64 predicted by the mean field approximation. This discrepancy is simply due to the fact that $N = 10$ is far from infinity. Yet, as we show below, the refined mean field allows one to obtain an approximation that is much more accurate.

To that, we replace the mean field expected queue length $\sum_{i \geq 1} \pi_i$ by its refined version $\sum_{i \geq 1} \pi_i + \frac{1}{N} \sum_{i \geq 1} V_i$ where V_i is the constant of Theorem 3. We report the values in Table 3.3 where we compare the refined mean field approximation with expected queue length that have been computed by simulation.

N	10	20	50	100	$+\infty$
Simulation (two-choice)	2.8040	2.5665	2.4344	2.3931	–
Refined mean field	2.7513	2.5520	2.4324	2.3925	2.3527
Simulation (pull-push)	2.3043	1.9700	1.7681	1.7023	–
Refined mean field	2.2945	1.9654	1.7680	1.7022	1.6364

Table 3.3: Expected queue length as a function of the system size N . We compare the refined mean field with the values that having computed by simulation. We choose $\rho = 0.90$ and $r = 1/(1 - \rho)$ for the pull-push.

This table illustrates that the refined approximation is much more accurate than the classical mean field approximation. While this table illustrates the accuracy with only two examples and one load ρ , we have shown that the good accuracy of the refined mean field by using other load balancing or coupon replication model in [P23] or simple epidemic examples in [P20].

3.4.4 Impact of Choosing With or Without Replacement

In addition to being more accurate, the refined approximation allows us to quantify the impact of phenomena that are indistinguishable by the classical mean field approximation. As an illustration, let us zoom on the two-choice model. When an incoming job arrives, two servers are sampled. In the literature, authors rarely distinguish if the two servers must be distinct or not, or said otherwise: if the two servers are sampled with or without replacement.

The reason for doing the distinction is that the two models coincide as N goes to infinity and are therefore indistinguishable by a mean field approximation. Indeed, the probability that least loaded of two randomly chosen servers has $i - 1$ jobs is $X_{i-1}^2 - X_i^2$ if the two servers are picked with replacement but is equal to $(X_{i-1} \frac{NX_{i-1}-1}{N-1} - X_i \frac{NX_i-1}{N-1})$ if the two servers must be distinct. As N goes to infinity, the two cases differ by a $1/N$ -term.

As shown in [P23], Theorem 3 can be refined to show that the expected queue length is equal to

$$\sum_{i \geq 1} \mathbb{E} [X_i^{(N)}] = \sum_{i \geq 1} \pi_i + \frac{1}{N} \sum_{i \geq 1} V_i - \frac{1}{N} \sum_{i \geq 1} E_i + O\left(\frac{1}{N^2}\right), \quad (3.17)$$

where V_i is the refined constant of the classical two-choice (with replacement) and E is equal to :

$$E_i = \sum_{j=1}^{\infty} \sum_{k=1}^{i-1} (1 - \pi_j) 2^{i-j-1} \rho^{2^i-2^j} = \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} (\rho^{2^{i+j}-2^j} - \rho^{2^{i+j}-1}) 2^{i-1}$$

	$N = 10$	$N = 20$	$N = 50$
with replac.	2.820 (rmf=2.751)	2.574 (rmf=2.552)	2.431 (rmf=2.432)
without replac.	2.705 (rmf=2.630)	2.502 (rmf=2.491)	2.406 (rmf=2.408)
with-without	0.115 (rmf=0.121)	0.073 (rmf=0.061)	0.026 (rmf=0.024)

Table 3.4: Two-choice : comparison with and without replacement. The values in parenthesis (rmf=·) correspond to the refined mean field approximation, the others to the values obtained by simulation. The rows “with-without” are the difference between “with replacement” and “without replacement”.

d	ρ	Mean field	$1/N$ -expansion	$1/N^2$ -expansion	Simulation
2	0.9	2.3527	2.7513	2.8045	2.8003
2	0.95	3.2139	4.1017	4.3265	4.2993
3	0.9	1.8251	2.2364	2.3322	2.3143
3	0.95	2.4084	3.3090	3.7604	3.5996

Table 3.5: d -choice model. Steady-state average queue length : comparison of the value computed by simulation with the three approximations (mean field and refined mean field of order $1/N$ and $1/N^2$). The system is composed of $N = 10$ servers.

As an illustration, we compare in Table 3.4 the refined approximation and values obtained by simulation for various values of $\rho = 0.90$ and different values of N . We observe that in most cases, the error E/N is a good predictor of the impact of choosing with or without replacement. We also observe that the value E/N tends to be smaller than the error term of the refined approximation V/N . This also explains why the impact of choosing with or without replacement has been overlooked in the literature: its impact is smaller than the error made by considering a mean field approximation.

3.4.5 Second Order Expansion

Theorem 3 can be extended to derive a second order expansion term (Theorem 4). For the steady-state estimation, this shows that

$$\mathbb{E} [X_i^{(N)}] = \pi_i + \frac{1}{N} V_i + \frac{1}{N^2} A_i + O\left(\frac{1}{N^3}\right).$$

Our conclusion from [P20] is that while adding the $1/N^2$ tends to improve the accuracy for many models, it is not always the case. In general, most of the gain in terms of accuracy are brought by the $1/N$ -term. As the $1/N^2$ -term is much more expensive to compute than the $1/N$ term, we believe that when the $1/N^2$ -expansion is too hard to compute, staying with the $1/N$ -expansion is already sufficient for many models.

As an illustration we compare in Table 3.5 the accuracy of three approximation (mean field, refined mean field with a $1/N$ term and refined mean field with also the $1/N^2$ term). In the case of the two-choice model ($d = 2$), adding a third order expansion seems to be beneficial and provides a more accurate estimate. On the contrary, for $d = 3$ and $\rho = 0.95$, the $1/N^2$ term seems to over-correct and provides an estimate that is not clearly better than the classical $1/N$ expansion. Note that this table only report values for $N = 10$. For larger values of N , the $1/N$ and $1/N^2$ expansion quickly coincide.

It can be shown that the fixed point of the two-choice model satisfies $\pi_i = \rho^{2^{i-1}}$ and that therefore the mean field expected queue length

$\sum_{i \geq 1} \pi_i$ is of order $-\log_2(1 - \rho)$. Based on our extensive simulations, we conjecture that the first order correction term of the refined mean field approximation is approximately equal to $\sum_{i \geq 1} V_i \approx \rho^2 / (2(1 - \rho))$ and that the second order approximation is $\sum_{i \geq 1} A_i \approx 1 / (20(1 - \rho)^2)$. This suggests that the average queue length of the two choice model is approximately equal to $a(N, \rho)$:

$$a(N, \rho) = \sum_{i=1}^{\infty} 2^{\rho^{i-1}} + \frac{\rho^2}{2N(1 - \rho)} + \frac{1}{20N^2(1 - \rho)^2} \quad (3.18)$$

While the first term is of order $\log \frac{1}{1 - \rho}$, the second term is of order $1 / (1 - \rho)$ when N is not too large. Hence, the above formula suggests that for a fixed N , the mean field approximation becomes less and less accurate as ρ approaches 1. This is what is observed in practice.

3.4.6 Transient Analysis

Theorem 3 shows that the refined mean field can be also be applied to study the transient regime of such a system. To illustrate this, we study how the expected queue length evolves with time in the two-choice model. The expected queue length can be expressed as $\sum_{i \geq 1} \mathbb{E}[X_i^{(N)}(t)]$. We choose the load $\rho = 0.9$ or $\rho = 0.95$. We choose an initial condition that depends on the load: for $\rho = 0.90$, out of the N queues, $0.2N$ queues start with 2 jobs and $0.8N$ queues start with 3 jobs; For $\rho = 0.95$, $0.5N$ queues start with 3 jobs and $0.5N$ queues start with 4 jobs. We choose this value as it is close to refined steady-state values (2.75 for $\rho = 0.9$ and 4.1 for $\rho = 0.95$).

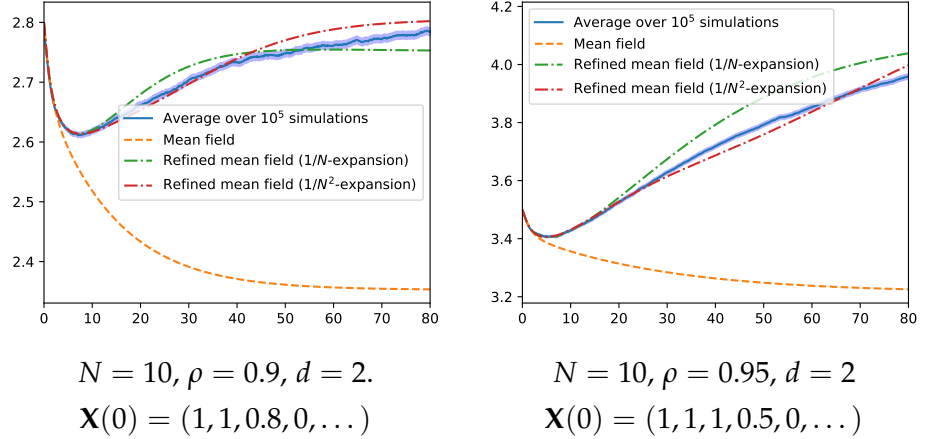


Figure 3.1: Two-choice model and transient regime: Comparison of the classical mean field approximation and the two expansions with data from simulations.

In Figure 3.1, we report how the expected queue length evolve with time compared to the three approximation (mean field, $1/N$ -approximation and $1/N^2$ -approximation). The blue region indicate

a 95% confidence interval. We observe in this figure, the expansions provide an estimation of the evolution of the expected queue length that is much more accurate than the one provided by the classical mean field approximation. The $1/N^2$ -expansion provides a better approximation than the $1/N$ -expansion. For larger values like $N \geq 20$, the expansion $1/N$ and $1/N^2$ are almost indistinguishable.

For the simulation of the transient regime, the running time of obtain one trajectory of simulation by using an optimize C++ simulator is approximately 0.05sec for $N = 10$ (and it grows linearly with N). The results represented in Figure 3.1 are averages over 10^5 simulations which represents roughly 1h of computation for each of the two panels. As a comparison, the total time to compute the expansion of order $1/N^2$ is about 10 seconds (and does not depend on N), and the time to compute the expansion of order $1/N$ is less than 1 second (using our python's implementation).

3.5 EXTENSIONS AND OPEN QUESTION

In this chapter, we present our results on two refinements of the mean field approximation. These approximations consists in adding a term in $1/N$ and potentially a second term $1/N^2$ where N is the size of the system. This result can be applied to a large class of models and are valid for the transient and stationary regime. Our examples show that the refined mean field approximation is much more accurate than the classical mean field approximation, especially for moderate system sizes (say $N = 10$ to $N = 100$).

The accuracy of the refined approximation leads to think that this refined approximation has potential applications in many domains. For some models, our results can be readily applied, especially by using our numerical implementation in [P18]. For some others, there remains a number challenges to solve and in particular regarding the **extension to other population models**.

- **Synchronous population model** – The extension of our results to the model that we described in Section 2.5 is quite natural. In [P22], we have developed a refined mean field approximation for such a model. The refined constant $V(t)$ is described by a linear recurrence equation instead of an ordinary differential equation. Our method to obtain this result is similar and also uses the computation of moments.
- **Multi-scale systems** – Multiple time-scale typically arise when a resource shared by everyone changes state at a faster evolution than each agent. This occurs for instance when a centralized controller interacts with a mass of agents by sending distributed control signals. The classical model of [8] can model such systems by considering a time-averaging method. Given the impor-

tance of such models in communication protocols, extending is an important open question.

- **Heterogeneous systems** – Many systems are non-uniform. This is the case for the example of caching example that we described in Section 2.1. For now, we do not have theoretical ground to explain why a refined mean field approximation would work in such a system. Yet, we will show in Chapter 4 that it is possible to derive such an approximation for this system that improves the accuracy compared to the classical mean field model.
- **Non-Markovian dynamics** – A typical application is the study of load balancing systems with non-exponential service time. Such systems have been for instance studied in [16], [1] by using a PDE approach for the transient regime or a method called the *queue at the cavity* for the stationary regime. We believe that the moment closure approach could be used to propose a *refined cavity* method.

Finally, one could ask whether it is worth looking at higher order expansion terms in $1/N^3$ or more. From a theoretical point of view, our method of moments could be directly extended to high moments, at the price of probably extremely complex derivations. From a practical point of view, when we compare the accuracy of the classical mean field approximation to the one of the expansions of order $1/N$ and $1/N^2$, most of the gain in terms of accuracy are brought by the $1/N$ -term. Hence, we doubt that higher order expansions can have a real practical interest.

HETEROGENEOUS MODELS AND CACHE ALGORITHMS

In this chapter, we study a class of cache replacement policies, called $\text{RAND}(\mathbf{m})$, $\text{LRU}(\mathbf{m})$. These algorithms adapt dynamically the content of the cache to the observed popularity of items without the need of learning it. One originality of this model is that it is a heterogeneous population model.

Most of the results on mean field approximation assume an homogeneous population of agents. A natural extension of such a model is to consider that the population of agents can be clustered into a finite number of classes with a large number of agents per class. This is what we did in Section 2.1. Yet, this model of a finite number of classes is often unrealistic and this is especially true for caching models where items are often assumed to have a Zipf popularity. In this chapter, we develop results on heterogeneous mean field that show that it is not necessary to use clustering to apply mean field approximation.

ROADMAP The chapter is organized as follows. After presenting the model in Section 4.1, we present an exact steady-state analysis in Section 4.2. This exact analysis allows us to disprove conjectures that date from the 80s. We then present a heterogeneous mean field approximation in Section 4.3 as well as a (partial) proof of its accuracy based on stochastic approximation. We then show how to compute a refined approximation in Section 4.4. Finally, we present some extensions and future work in Section 4.5.

The results contained in this chapter come from [P26] for the first analysis, from the recently submitted paper [P3] for the refined approximation and for [P25] for the extensions to LRU. It also covers some of the materials of [P15], [P28], [P27].

4.1 THE $\text{RAND}(\mathbf{m})$ POLICIES AND VARIANTS

In this section we introduce a family of replacement algorithms, called $\text{RAND}(\mathbf{m})$. The $\text{RAND}(\mathbf{m})$ policy is a generalization of the RANDOM policy introduced in Section 2.1. It makes use of h lists. $\mathbf{m} = (m_1, m_2, \dots, m_h)$ is a tuple of integers where $m_i \geq 1$ denotes the size of the i th list. Items enter the set of lists via List 1 and whenever requested while being part of List i they move up one list. More specifically, one of the following three events occurs when an item, say item k , is requested at some point in time:

1. If Item k is not in any of the lists (miss) – in this case item k is inserted in List 1 by replacing a random item from this list.
2. If Item k is in List $i < h$ (hit) – in this case, it is exchanged with an item from List $i + 1$ (picked at random).
3. If Item k is in List h (hit) – no changes are made.

Note: the behavior described above implicitly assume that the cache is always full. When the cache contains fewer elements than its capacity, we consider that the remaining capacity is filled with “dummy” items that will progressively leave the cache as they are replaced by “real” items.

As in Section 2.1, we assume the IRM model: the requests of Item k arrive according to a Poisson process of intensity λ_k . Without loss of generality, we assume that $\sum_{k=1}^n \lambda_k = 1$ so that λ_k is also the probability that a given request is for item k . All items have the same size and the cache has size $m = \sum_{i=1}^h m_i$. Note that two classical policies can be represented by this model:

- RANDOM – The case with a single list $\mathbf{m} = (m)$ corresponds to the RANDOM policy.
- CLIMB – The case where all lists have size 1, $\mathbf{m} = (1, \dots, 1)$, is called the CLIMB algorithm or the TRANSPOSITION rule in the literature [79], [44].

The RAND(\mathbf{m}) policy can be obtained by slightly modifying a family of replacement algorithms introduced in [3]. We also studied variants of this policy where the RANDOM exchanges are replaced by FIFO or LRU. This leads to the policies FIFO(\mathbf{m}) in [P26] or LRU(\mathbf{m}) in [P25].

4.2 EXACT STEADY-STATE ANALYSIS

4.2.1 The Product-form Stationary Measure

Let $Y_i(t)$ be the set of items that are in the i th list at time t (with $1 \leq i \leq h$). It should be clear that, under the RAND(\mathbf{m}) policy, \mathbf{Y} is a continuous time Markov chain that evolves in a state space $\mathcal{Y}(\mathbf{m})$. Our first result [P26, Theorem 1] is an exact characterization of the steady-state distribution of \mathbf{Y} . We show that it satisfies a product-form:

An element of $\mathcal{Y}(\mathbf{m})$ is a tuple $\mathbf{y} = (y_1 \dots y_h)$ such that y_i is a subset of $\{1 \dots N\}$ that contains exactly m_i element where all the subsets y_i are disjoint.

$$\pi(\mathbf{y}) = \frac{1}{Z(\mathbf{m})} \prod_{i=1}^h \sum_{k \in y_i} (\lambda_k)^i, \tag{4.1}$$

where $Z(\mathbf{m})$ is a normalization constant: $Z(\mathbf{m}) = \sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{m})} \prod_{i=1}^h \sum_{k \in y_i} \lambda_k^i$.

The proof of this result is quite direct and uses the reversibility of the process \mathbf{Y} . The above formula notably generalizes the formula of CLIMB and RANDOM of [35], [53].

4.2.2 Hit Probability and Numerical Algorithms

The steady-state hit probability is the probability than when an item is requested, this item is found in the cache. By using Equation (4.1), it can be written

$$\text{Hit}(\mathbf{m}) = \sum_{\mathbf{y} \in \mathcal{Y}(\mathbf{m})} \pi(\mathbf{y}) \sum_{i=1}^h \sum_{k \in y_i} \lambda_k. \quad (4.2)$$

While plugging Equation (4.1) into this formula could be called a closed form expression, a direct computation of the hit probability using this formula is prohibitive because the number of configuration $\mathcal{Y}(\mathbf{m})$ is extremely large.

We show in [P26, Section 4.2] this computation is possible by using dynamic programming. The idea is that the hit probability of a population of N items and lists sizes \mathbf{m} lists can be expressed recursively as a function of the hit probabilities of the systems with $N - 1$ items and list sizes $\mathbf{m} - \mathbf{e}_i$. The derivation of the recursive formula is rather easy once the good notations are chosen. The main difficult for implementing it is to avoid underflow. More details can be found in [P26, Section 4.2].

4.2.3 CLIMB is not Optimal

By using Equation (4.2) and a variant of the FKG inequality [32] that can be found in [10], we show that among all $\text{RAND}(\mathbf{m})$ variants that use a cache of size m , having a single list $\mathbf{m} = (m)$ always leads to the lowest hit probability. Said otherwise, $\text{RAND}(\mathbf{m})$ always outperform RANDOM . It is tempting to conjecture that separating a list into two smaller lists always improves the performance and that the CLIMB algorithm, that is, having m lists of size 1, achieves the lowest miss rate within this class. In [3, p135] an even stronger conjecture is presented that states that CLIMB is optimal under the IRM model for all finite-memory demand replacement algorithms.

policy	\mathbf{m}	Miss probability
Optimal	$\text{RAND}(1,1,4)$	0.005284
	$\text{RAND}(1,1,3,1)$	0.005299
	$\text{RAND}(1,1,1,3)$	0.005338
CLIMB	$\text{RAND}(1,1,1,1,1,1)$	0.005348
LRU	$\text{LRU}(6)$	0.005880
RANDOM	$\text{RAND}(6)$	0.015350

Table 4.1: Counter-example that shows that CLIMB is not optimal for the IRM model: $p = (49, 49, 49, 49, 7, 1, 1) / 205$ and $m = 6$.

In Equation (4.2), $\sum_{i=1}^h \sum_{k \in y_i} \lambda_k$ is the probability that a request item is in the cache when the cache state is \mathbf{y} .

We show in [P26] that this conjecture is false by exhibiting a probability distribution for which there are seven configurations for $\text{RAND}(\mathbf{m})$ that outperform CLIMB. We show some of the best performing distribution in Table 4.1. For comparison, we also show the performance of RANDOM and LRU. This shows that CLIMB is not optimal. It further indicates that when fixing the number of lists h , the optimal choice for the length of each list is not necessarily setting $m_1 = \dots = m_{h-1} = 1$: for instance, $\mathbf{m} = (1, 1, 3, 1)$ is better than $\mathbf{m} = (1, 1, 1, 3)$. This demonstrates that another of the *natural* conjectures formulated in [3, p135] is also false.

4.3 TRANSIENT ANALYSIS VIA MEAN FIELD

The previous section shows that it is possible to use an exact analysis to characterize the steady-state performance of a $\text{RAND}(\mathbf{m})$ policy. Yet, an exact analysis suffers from two main drawbacks: First it is only applicable to steady-state performance and provides no clue on the transient regime and in particular how the parameter \mathbf{m} affects the rate at which the replacement policies adapt to changing popularities. Second, our recursive formula to compute Equation (4.2) is polynomial in the number of items but exponential in the number of lists.

In this section we show how to derive a mean field approximation for the $\text{RAND}(\mathbf{m})$ policy. The originality of this mean field approximation is that the system is **heterogeneous** because all items have different popularity. This approximation becomes exact as the number of items and the cache size tends to infinity and allows for a fast and accurate study the transient behavior of the algorithm. We illustrate how this approximation can be used to compute the time to fill an empty cache or obtain a fast approximation of the steady-state miss probability.

4.3.1 Derivation of the Equations and Intuition

At a given time step t , Item k is either part of some list $i \in \{1, \dots, h\}$ or is not in the cache. If an item is not in the cache, we say that it is part of List 0. For an item $k \in \{1 \dots n\}$ and a List $i \in \{0 \dots h\}$, we define the random variables $X_{k,i}(t)$, where $X_{k,i}(t)$ equals 1 if the Item k is part of List i at time t and 0 otherwise. The probability that Item k is in List i at time t is $\mathbb{E}[X_{k,i}(t)]$.

Item k is requested at rate λ_k and thus moves from list i to List $i + 1$ at rate λ_k . The transitions for this model are (for $i \in \{0 \dots h - 1\}$ and any items k, ℓ):

$$\mathbf{X} \mapsto \mathbf{X} - \mathbf{e}_{k,i} + \mathbf{e}_{k,i+1} + \mathbf{e}_{\ell,i} - \mathbf{e}_{\ell,i+1} \quad \text{at rate } \lambda_k X_{k,i} \frac{X_{\ell,i+1}}{m_{i+1}} \quad (4.3)$$

In [P3], we consider a slightly more general variant where λ_k can also depend on the List i .

As in Section 2.1, we construct an approximation of $\mathbb{E}[X_{k,i}(t)]$ by a deterministic quantity $x_{k,i}(t)$. The initial conditions of the ODE are $x_{k,i}(0) = 1$ if the Item k is in the i th list of the cache at time $t = 0$ and 0 otherwise. Let $(x_{k,i}(t))_{k,i}$ be the unique solution of the following set of ODEs, for $k \in \{1 \dots n\}, i \in \{1 \dots h\}$:

$$\begin{aligned} \dot{x}_{k,i}(t) = & \lambda_k x_{k,i-1}(t) - \sum_{\ell} \lambda_{\ell} x_{\ell,i-1}(t) \frac{x_{k,i}(t)}{m_i} \\ & + \mathbf{1}_{\{i < h\}} \left(\sum_{\ell} \lambda_{\ell} x_{\ell,i}(t) \frac{x_{k,i+1}(t)}{m_{i+1}} - \lambda_k x_{k,i}(t) \right), \end{aligned} \quad (4.4)$$

where $\mathbf{1}_{\{i < h\}}$ is equal to 1 if $i < h$ and 0 otherwise.




Figure 4.1: Evolution of the list in which Item k is. When λ_k is small and the m_i s are large, the state of one item becomes independent of the hit rate in each box. Its behavior can be approximated by a time-inhomogeneous continuous-time Markov chain. This is the mean field approximation.

This equation can be understood as follows. Assume that Item k is in List $i \in \{0 \dots h-1\}$ at time t . At rate λ_k , Item k is requested and moves to List $i+1$. At rate $H_{i-1}(t) = \sum_{\ell} \lambda_{\ell} X_{\ell,i-1}(t)$, an item from List $i-1$ is requested and is exchanged with an item from list i chosen at random. This item is Item k with probability $1/m_i$. Hence, with probability $H_{i-1}(t)/m_i$, Item k moves to List $i-1$. If the list in which Item k is and the variables $H_i(t)$ were independent, the behavior of Item k would be described by a Markov chain whose transition matrix is represented in Figure 4.1.

If the quantities $H_i(t)$ were deterministic, the Markov chain represented in Figure 4.1 is a birth-death process. Hence, the stationary measure of this chain can be easily computed: the probability that this chain is in state i in steady-state is $\pi_{k,i} \propto (\lambda_k)^i \prod_{j=0}^{i-1} \frac{m_j}{H_j}$. Denoting $z_i = \prod_{j=0}^{i-1} \frac{m_j}{H_j}$, this shows that a fixed point of the ODE satisfies:

$$\pi_{k,i} = \frac{p_k^i z_i}{1 + \sum_{j=1}^h p_k^j z_j}.$$

We show in [P26, Theorem 7] that the ODE (4.4) has a unique fixed point that satisfies the above equation and the cache size constraint $\sum_{i=1}^n \pi_{k,i} = m_i$. By using a Lyapunov function based on relative entropy, we also show in [P15] that this fixed point is an attractor. This fixed generalizes the Equation (2.2) and can be computed efficiently by an iterative scheme.

It should be clear that this ODE is Lipschitz-continuous and hence has a unique solution. Moreover, with the initial conditions indicated, the solution of the ODE satisfies that, for all time t : $\sum_{k=1}^n x_{k,i}(t) = m_i$ and $\sum_{i=0}^h x_{k,i}(t) = 1$.

α	n	m_1	m_2	exact	mean field
0.8	300	2	98	0.3466	0.3470
		98	2	0.4239	0.4245
0.8	3000	20	980	0.3034	0.3035
		980	20	0.3723	0.3724

Table 4.2: Mean field model validation for $h = 2$ with Zipf-like popularity distributions.

4.3.2 Accuracy of the Mean Field Approximation

In [P26], we give two arguments to justify the accuracy of the approximation. The first is a theoretical result [P26, Theorem 6] that shows that, when the number of items goes to infinity and cache size is large, the popularity of the different lists are in the cache converge to their mean field approximation. More precisely, by denoting by $H_i(t)$ the popularity of the items that in List i : $H_i(t) = \sum_k \lambda_k X_{k,i}(t)$ and $h_i(t) = \sum_k \lambda_k x_{k,i}(t)$ its mean field approximation, we show that for any time horizon T , one has

$$\mathbb{E} \left[\sup_{t < \frac{T}{\max_k \lambda_k + \max_i 1/m_i}} |H_i(t) - h_i(t)| \right] \leq C \sqrt{\max_k \frac{\lambda_k}{\sum_i \lambda_i} + \max_i 1/m_i}. \quad (4.5)$$

This result is a sample path result similar that we obtained by stochastic approximation argument as in Section 2.3. This result is quite strong and (to the best of our knowledge), it is one of the few results that look at completely heterogeneous models and that does not cluster the objects into a finite number of classes with a high number of objects per class. We have extended this results to LRU-based caching policies in [P25] by essentially showing that the time that an object spends in the cache becomes deterministic. Note that our approach differs from the one [33], [49] that heavily rely on a precise characterization of the stationary measure of the LRU policy in terms of TTL.

Our second argument to justify the accuracy of the approximation is based on a numerical study. As an illustration, we compare in Table 4.2 the steady-state miss probability given by the mean field approximation with the exact values. The popularity of the item follows a Zipf-like distribution and we vary \mathbf{m} . We observe that the mean field approximation is within 1% for all cases and seems smaller for larger cache sizes. As an other illustration, we depict in Figure 4.2 how the popularity of the cache evolves with time. The system is composed of 1000 items that follows a Zipf popularity. The cache can contain up to 200 items and is initially empty. We compare a case with

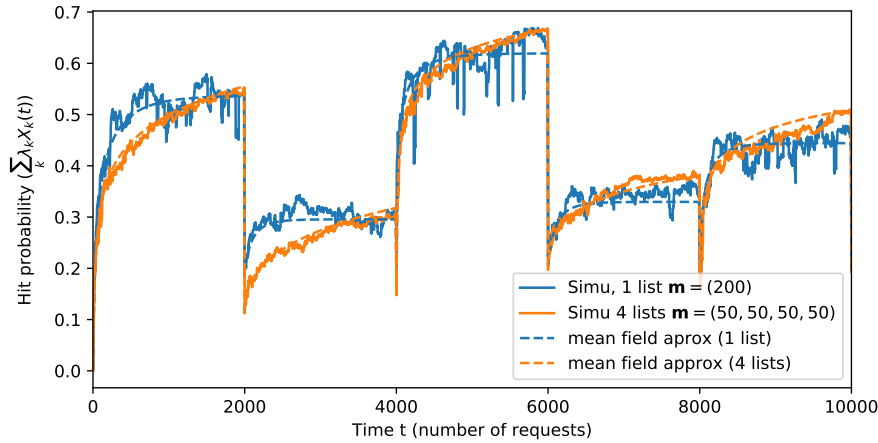


Figure 4.2: Cache and transient regime

only one list (pure RANDOM policy or $\text{RAND}(200)$) and a case with four lists of equal size (*i.e.* $\text{RAND}(50,50,50,50)$). In order to illustrate how the two policies adapt to changing popularities, we randomly swap the popularities of the various items in the cache every 2000 requests by picking a new parameter α between uniformly between 0.50 and 1 and randomly permuting the popularities of all item so that item 1 is not always the most popular. This figure shows an excellent agreement between the mean field (ODE) approximation and the simulation: the mean field approximation is essentially a denoised version of the simulation. This figure shows that, as expected, $\text{RAND}(50,50,50,50)$ has a higher steady-state hit rate than RANDOM but takes more time to adapt to changing popularities.

These results make us believe that neither our accuracy results nor the ones of [33], [49] really explain why the mean field approximation ODE (4.4) or the TTL approximation studied in [33], [49] perform that well. The reason for that is that our method to obtain the bound of Equation (4.5) is based on stochastic approximation and almost sure convergence. We believe that using a generator approach such as the one we developed in Chapter 3 would lead to much tighter bound. In the next section, we show in fact that the refined framework can be used to construct a more accurate approximation, even if we do not have a proof of its accuracy for now.

4.4 A HETEROGENEOUS REFINED APPROXIMATION

4.4.1 Construction of the Refined Approximation

It is possible to transform the model described in Equations (4.3) into a density dependent population process by scaling the size of the

transitions by a factor $1/N$ and accelerating the time by a factor N . For a given scaling factor N , the transitions of the model are:

$$\mathbf{X}^{(N)} \mapsto \mathbf{X}^{(N)} + \frac{1}{N}(-e_{k,i} + e_{k,i+1} + e_{\ell,i} - e_{\ell,i+1}) \quad \text{at rate } N\lambda_k X_{k,i} \frac{X_{\ell,i+1}}{m_{i+1}} \quad (4.6)$$

The case $N = 1$ corresponds to the original model.

This allows us to use the result of Theorem 3 in order to define and compute a refined mean field approximation of $\mathbb{E}[\mathbf{X}_{k,\ell}^{(N)}(t)]$ as follows:

$$\mathbb{E}[\mathbf{X}_{k,i}^{(N)}(t)] = x_{k,i}(t) + \frac{1}{N}V_{k,i}(t) + O\left(\frac{1}{N^2}\right),$$

where $x_{k,i}(t)$ is the classical mean field approximation and V is the solution of a linear ODE.

In order to obtain a refined approximation for the original model, we propose to apply the above formula with $N = 1$. This leads us to define a refined quantity $y_{k,i}(t)$ as follows:

$$y_{k,i}(t) = x_{k,i}(t) + V_{k,i}(t) \quad (4.7)$$

Note that Theorem 3 guarantees that the refined mean field in theory accurate as N goes to infinity. Here, we apply this formula with $N = 1$ that can hardly be considered as close to infinity. Yet, as will see below, this refined approximation is remarkably accurate.

4.4.2 Empirical Validation

To assess how accurate is the refined approximation Equation (4.7) compared to the classical mean field, we used our tool [P18] to perform a numerical study. We consider two cases ($n = 10$ and $n = 100$ items). The popularity of the items follows a pareto-distribution of parameter $\alpha \in \{0.6, 1, 1.4\}$: $\lambda_i \propto i^{-\alpha}$. we vary the cache size $m \in \{0.1n, 0.25n, 0.5n\}$ and number of lists between $h \in \{1, 2, 5\}$. The list have all equal sizes: $\lceil mn/h \rceil$.

To compare the two approximation, we estimated the miss probability by using the two approximations and compare this to data obtained by simulations. For each of the two approximations, compare two metrics:

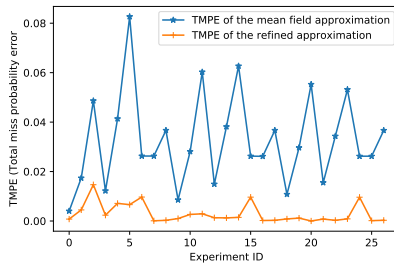
- (a) The total miss probability error (TMPE);

$$TMPE = \left| 1 - \frac{\text{total miss proba (computed by approximation)}}{\text{total miss proba (computed by simulation)}} \right|$$

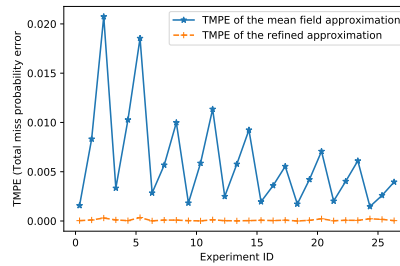
- (b) The average per-item miss probability error (IMPE):

$$IMPE = \left| 1 - \frac{\text{miss proba of Item } k \text{ (by approximation)}}{\text{miss proba of Item } k \text{ (by simulation)}} \right|$$

In Figure 4.3, we report the total miss error probability as a function of the experiment ID. We compare a case with $n = 10$ (left panel) items and a case with $n = 100$ items (right panel). There are various observations than can be made. First, the two approximations are an order of magnitude more accurate for $n = 100$ compared to $n = 10$. Second, the refined approximation is significantly more accurate than the classical mean field approximation, especially for $n = 100$ items. Last, the error of the mean field approximation seem to decrease with the exponent α . This last fact can be explained by the fact that the total miss probability is small in such a case which in turn means that the denominator in the definition of the TMPE is smaller for larger α .



$n = 10$ items



$n = 100$ items.

Figure 4.3: Total miss probability error (TMPE) as a function of the experiment ID.


Note that when averaging over all experiments, we obtain the average errors reported in Table 4.3. Note that the error of the refined mean field is approximately 10 times smaller than the one of the classical mean field for $n = 10$ items and and 300 times smaller for $n = 100$ items.

N	Mean field TMPE	Refined mean field TMPE
10	0.032759	0.002797
100	0.006081	0.000019

Table 4.3: Average TMPE of both approximation

In Figure 4.4, we report the per item miss error probability (IMPE) as a function of the experiment ID. As for the TMPE the refined approximation is more accurate than the classical mean field approximation, although the improvement is less impressive. Moreover, measured in terms of IMPE, the accuracy of both approximations seems to diminish with the number of lists h and the exponent α . In fact, when both h and α are large, some items have a miss probability close to 0, which lead the average IMPE to be large if these items are not perfectly estimated. This is also related to what we observed in [P23]: while the refined mean field approximation estimate very accurately the shape

There are 27 experiments first 9. The parameters as a function of the experiment ID are as follows:



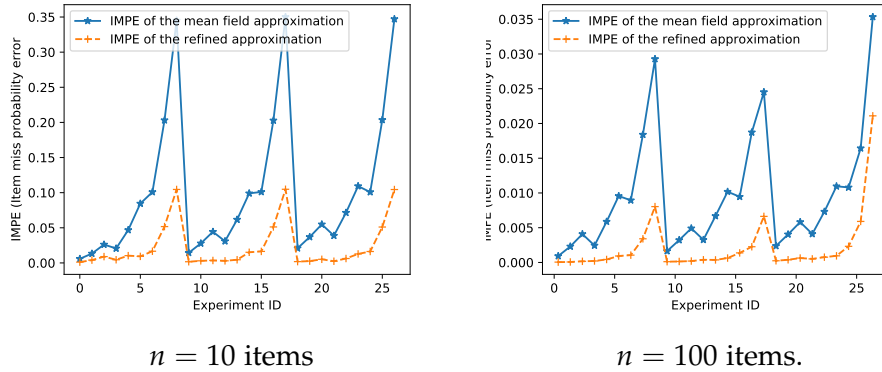


Figure 4.4: Item miss probability error (IMPE) as a function of the experiment ID.

of the distribution, it does not estimate very accurately the tail of the distributions.

In [21], another approximation has been proposed that works by using an asymptotic expansion of the normalization constant $Z(\mathbf{m})$ defined in Equation 4.1. Measured in terms of TMPE, this approximation performs similarly than our refined approximation: it is more accurate on some examples and less on others. Measured in terms of IMPE, this approximation is significantly more accurate because it estimates more precisely the miss probability of the items that have a very small miss probability. Note that the technique used to obtain this other approximation relies on having a product form stationary measure and is quite specific to this example.

4.5 EXTENSIONS AND OPEN QUESTIONS

In this chapter, we studied a model of cache replacement policies. One of our most original contribution to this subject is to manage to define a heterogeneous mean field approximation and to prove that it is asymptotically accurate. Note that after this work, we also extended the results to the LRU-based policies LRU(\mathbf{m}) and h -LRU in [P25], [P28] in which we study a TTL-approximation of those policies. In these papers, we also derive a proof of the accuracy of those TTL-approximation by using a similar stochastic approximation technique.

A second contribution of our work is to have shown that it is possible to use our refined mean field approximation for a completely heterogeneous system, by using this approximation with a scaling parameter $N = 1$. Yet, the application of this refined approximation raises two main challenges. First, we have for now no theoretical justification of the approach. For future work it seems fundamental to understand to which heterogeneous models this approximations can be applied. Second, the numerical complexity is problematic. By using

our current framework, computing a refined mean field for a model of cache with n items and h lists requires solving a linear system with $(nh)^2$ variables which makes the computation of this approximation impossible for more than 100 or 200 items.

SPATIAL INTERACTIONS AND BIKE-SHARING SYSTEMS

Bike-sharing systems are becoming important for urban transportation. In such systems, users arrive at a station, take a bike and use it for a while, then return it to another station of their choice. In this chapter, we study a model of bike-sharing systems. We use this model to study dimensioning problems, rebalancing algorithms and forecasting.

One particularity of bike-sharing system is that the geometry takes an important role. For instance, when a user arrives at a station that is full, it will most probably look at neighboring station. This makes the mean field approximation inapplicable in theory. We show that in practice, mean field approximation can still be applied as is and gives very reasonable result. We also show how the original mean field model can be refined to take into account spatial interactions by using a method called *pair approximation*.

ROADMAP This chapter is organized as follows. We first describe how a bike-sharing model can be viewed as a close queuing network in Section 5.1. We develop a mean field approximation for such a system (§5.1.1) which allows us to study simple incentives (§5.1.2). While this mean field approximation does not take into account geometry, we show by simulation that this mean model is close to a model that would take into account geometry (§5.1.3). The mean field models allows us to approximate the behavior of each station by a time-dependent M/M/1 queue. We then propose a model than is a refinement of the mean field approximation to take into account spatial interactions in Section 5.2. We show how the mean field approximation can be used to build a forecasting tool in Section 5.3. Finally, we conclude by some open questions in Section 5.4.

This chapter is mostly based on [P7] for the original bike-sharing model, on [P14] for the work on forecasting and on [P19] for the work about pair approximation.

5.1 THE BIKE SHARING MODEL

We consider a Markovian model of a bike-sharing system with N stations and a fleet of $\lfloor sN \rfloor$ bikes (*i.e.*, s bikes per stations). A bike can be either locked at a station or in transit between two stations. In this section, as in [P7], we mainly focus on the homogeneous bike-sharing model. This model enables us to obtain a closed-form expression for the optimal performance and to investigate incentives

and redistribution by trucks in this framework. This study is extended to an inhomogeneous model by [P6]. We also use the inhomogeneous model to study the forecasting problem in Section 5.3.

We assume that a station i can host K_i bikes. At a given station i , users arrive at rate $\lambda_i(t)$. If the station has one or more bikes, the user picks up a bike at this station and chooses a destination j with probability $p_j(t)$, otherwise she leaves the system. When the user arrives at a station j , if the destination has fewer than K_j bikes, the user returns her bike to this station. Otherwise, the user visits randomly other stations until she finds a station with fewer than K_j bikes. The trip time between the two stations is exponentially distributed with mean $1/\mu_{ij}$.

5.1.1 Homogeneous Mean Field Model

The homogeneous model is the simplest bike-sharing model. No quantity depends on time or on the station: $\lambda_i(t) = \lambda$, $p_j(t) = 1/n$, $\mu_{ij} = \mu$ and $K_i = K$. The advantage of such a model is that it can be thoroughly analyzed and leads to closed-form results. This model does not incorporate any geographical information.

Let us denote by $X_k^{(N)}(t)$ the proportion of stations that have k bikes available at time t . $\mathbf{X}^{(N)}$ is a Markov chain. There are two types of transitions. The first corresponds to a bike that is picked up. The second corresponds to a bike that arrives at a station. The number of bikes in transit is equal to sN (the total number of bikes) minus $N \sum_{k=1}^K ky_k$ (the number of bikes locked at the stations). Hence, the rate of arrival of bikes at a given station is $\mu(N - \sum_{k=1}^K ky_k)$. This shows that the transitions of the Markov chain are:

$$\begin{aligned} \mathbf{X}^{(N)} &\mapsto \mathbf{X}^{(N)} + \frac{1}{N}(\mathbf{e}_{k-1} - \mathbf{e}_k) \text{ at rate } \lambda N x_k \mathbf{1}_{k>0} \\ \mathbf{X}^{(N)} &\mapsto \mathbf{X}^{(N)} + \frac{1}{N}(\mathbf{e}_{k+1} - \mathbf{e}_k) \text{ at rate } N x_k \mu \left(s - \sum_{n=0}^K n x_n \right) \mathbf{1}_{k<K} \end{aligned}$$

where the k -th unit vector of \mathbb{R}^{K+1} is denoted by \mathbf{e}_k and $\mathbf{1}_{k<K}$ is equal to $\mathbf{1}$ when $k < K$ and 0 otherwise.

This defines a density dependent population process as defined in Section 2.2. This shows that, as N goes to infinity, $\mathbf{X}^{(N)}$ converges to the solution of a deterministic ODE. We show in [P7, Section 3.1] that this ODE has a unique fixed point to which all trajectories converge. The proof of this attractiveness uses a Lyapunov function constructed by using relative entropy.

This fixed point is of the form $x_i \propto \rho^i$ where ρ depends on the fleet size with $\rho < 1$ if $s < K/2 + \lambda/\mu$ and $\rho > 1$ if $s > K/2 + \lambda/\mu$. We say that a station is problematic if it has either 0 or K bikes. By using the mean field approximation, we show that the proportion of problematic stations depends on the fleet size s . It is minimal when $\rho = 1$, which

occurs when the number of bike per station s is slightly more than half of the number of places in the system: $s = K/2 + \lambda/\mu$. In this case the proportion of problematic stations is $2/(K + 1)$. This amounts to about 6.5% of problematic stations when the capacity of one station is $K = 30$ bikes.




Figure 5.1: Illustration of the steady-state occupancy of station when the capacity of stations is $K = 10$. The system is balanced when the number of bikes is equal to $K/2 + \lambda/\mu$.

In practice, most of the existing bike-sharing systems satisfy this rule of 50% of bikes compared to the number of slots. Note that we had the opportunity to discuss with the responsible of bike-sharing system installed in the city of Pisa that reduced their fleet size from 70% to 50% of bikes after our discussion. This reduced the congestion in their system.

5.1.2 Incentives and Regulation

In [P7], we consider two ways for improving the performance. The first is by what we call *incentives*. We model this by assuming that when a user wants to return her bike, she indicates two stations and the bike-sharing system indicates to her which one of the two has the least number of bikes available. We show that, when the two stations are picked *at random*, the proportion of problematic stations can be as low as $\sqrt{K}2^{-K/2}$ (instead of $2/(K + 1)$ in the original model without incentives). The performance is thus improved dramatically, and even if only a small percentage of users obey this rule. The proof of this result is again based on studying the fixed point of a density dependent population process.

We also study the problem of rebalancing the number of bikes by moving them by using trucks that move bikes from saturated stations to empty ones at rate γ . We show that if the bikes are moved one by one, the minimal redistribution rate needed to suppress any problematic station is equal to λ/K where K is the capacity of stations. Note that our simulations suggest that having larger truck size does not affect qualitatively the performance. This model can again be cast as a density dependent population process but with an additional

Such incentives are meant to encourage users to return a bike to a station that has few bikes. This is implemented successfully in New York city via the “bike angel” program [25].

difficulty: the drift is not continuous. This lead us to apply the results on discontinuous drift and differential inclusions that we developed in [P11]. This result is illustrated in Figure 5.2 in which we show that a large redistribution rate removes all problematic stations.




Figure 5.2: Illustration of the impact of redistribution. The station capacity is $K = 10$ and $\lambda = \mu = 1$. The fleet size is $s = 7$ and the redistribution rates varies from $\gamma = 0$ (no redistribution) to $\gamma = 1/5$.

Note that the variants that we discuss here focus on a time-homogeneous uniform model. The performance of an inhomogeneous model are of quite different nature and are studied in [P6].

5.1.3 Validation by Simulation

In [P7, Section 6], we validate the various hypothesis of the model by comparing our theoretical result with simulations of a model that violates one or the other hypothesis. Our simulations indicates that our results are quite robust with respect to this small changes of hypothesis. For instance, the trip time distribution seems to have very little effect on the performance (only the mean matters). We do not think that the model with N stations is insensitive to the trip time distribution. Yet, our result suggest that the system might be asymptotically insensitive when the number of station grows. This asymptotic insensitivity has been observed in other cases and we believe that similar proofs could be applied [86]: essentially, the arrival process at a station is asymptotically a Poisson process.

Among the other variants of the model, one of the most interesting is the influence of geometry on the performance metrics. Our theoretical results are based on a mean field model and therefore ignore the geometry of the system. In real system, a user that aims at a particular destination might accept to go to a neighboring station but will not go to the other side of city. Studying such a system analytically is out of reach of our method. In [P7, §6.3], we present simulation results that show that the influence of geometry on the proportion of problematic stations is quite limited, both for the basic model and the two-choice model. We show, however, that geometry has a much larger effect when other metrics are considered, such as the time needed to return a bike.

In the following section, we show how geometry can be better taken into account by using a technique called pair approximation.

5.2 SPATIAL REFINEMENT AND PAIR APPROXIMATION

Motivated by our application to bike-sharing systems, we study an abstract queuing model with geometry in [P19]. The model that we consider is a variant of the two-choice model that we described in Section 3.4. In the original model of [90], [67], when a task arrives, it is allocated to the least loaded of two servers picked at random among a collection of n servers. Here we study what happens when servers form a graph and that a task is routed on the least loaded among two neighbors. While the bike-sharing model was a closed queuing network, we study what happens to the more classical open queuing network in which jobs leave after having been served.

To the best of our knowledge (and as argued in [11]) there are very few studies of load balancing systems where geometry is taken into account. An exception is the recent work in [69] in which the authors consider a join-the-shortest queue policy in a graph with a large number of neighbors. The authors show that for such a system a mean field approximation is accurate, even if showing the accuracy is much harder than in the non-geometric model. Our focus here is different as we study the case where each server has few neighbors (typically 2 to 4).

As we will see, the mean field approximation of such a system is not very accurate. To obtain a more accurate approximation, we use a technique called pair-approximation, see for instance [72]. This technique shares some similarity with the moment closure technique that we used in Section 3.3 to develop our refined approximation.

5.2.1 Geometric Two-choice model

Our system is composed of N identical servers that are connected by an undirected graph (V, E) , where the set of vertices is the set of servers $V = \{1 \dots N\}$. Each server serves jobs at rate 1 and uses a first-come first-serve discipline (jobs sizes are exponentially distributed). Jobs arrive at each server at rate $\rho < 1$. When a job arrives at a server, say s_1 , another server s_2 is sampled uniformly at random among all neighbors of s_1 . The job is then allocated to the server s_1 or s_2 that has the least number of jobs (ties are broken at random). This allocation scheme is similar to the one of [52].

When the graph is complete, this model corresponds to the classical two-choice model. In this section, we will consider other variants:

- *Ring* – This model is illustrated in Figure 5.3a. Two servers s_1 and s_2 are neighbors if $s_1 = s_2 \pm 1$ (modulo N).
- *2D torus* – This corresponds to a 2D grid of size $\sqrt{N} \times \sqrt{N}$. A server in position (x, y) has four neighbors: $(x \pm 1, y)$ and $(x, y \pm 1)$ (modulo \sqrt{n}). This case is shown on Figure 5.3b.




Figure 5.3: Geometric two-choice models.

Note that the notation X_i used in Section 3.4 was different: in Section 3.4 X_i represents the proportion of servers having at least i jobs whereas here X_i is the proportion of servers having exactly i jobs.

- *Fixed degree* – We also simulate random graphs with fixed degree k . For each server, k neighbors are picked at random (for simplicity of generation, we allow self loop, like Node 5 of Figure 5.3c). The interaction graph remains constant during the simulation.

5.2.2 Mean Field Approximation

When the interaction graph is complete, all servers are exchangeable. For $i \in \{0, 1, 2, \dots\}$, let $X_i(t)$ be the proportion of servers that have i jobs at time t . $\mathbf{X} = (X_0, X_1, \dots)$ is a Markov chain whose transitions are as follows: there is a departure from a server with i jobs at rate $NX_i(t)$. When there is an arrival, the two chosen servers have i and j jobs with probability $X_i(t)X_j(t)$. If $i = j$ or $i < j$, the job is allocated on a server with i jobs. If $j < i$, the job is allocated on a server with j jobs. As a result, there is an arrival in a server with i jobs at rate $\rho(X_i(t))^2 + 2\rho \sum_{j=i+1}^{\infty} X_i(t)X_j(t) = 2\rho X_i(t)P_i(t)$, where $P_i(t) = X_i(t)/2 + \sum_{j=i+1}^{\infty} X_j(t)$ is the probability that an arrival on a server with i jobs is allocated to this server. This shows that \mathbf{X} is a density dependent population process whose mean field approximation is

$$\dot{x}_i = (x_{i+1} - x_i \mathbf{1}_{\{i>0\}}) + 2\rho(\mathbf{1}_{\{i>0\}} p_{i-1}(\mathbf{x})x_{i-1} - p_i(\mathbf{x})x_i), \quad (5.1)$$

where $p_i(\mathbf{x}) = x_i/2 + \sum_{j=i+1}^{\infty} x_j$.

5.2.3 The Pair Approximation Equations

We now consider a general interaction graph in which all nodes have the same degree k . Let $Y_{i,j}(t)$ the proportion of connected pairs of servers that have (i, j) jobs and $X_i(t) = \sum_j Y_{ij}(t)$ the proportion of servers that have i jobs. When the graph is complete, $Y_{ij}(t) = X_i(t)X_j(t)$, and \mathbf{X} is a Markov chain. This does not hold when the graph is not complete. A randomly chosen neighbor of a randomly chosen server having i jobs has j jobs with probability $Y_{ij}(t)/X_i(t)$. Hence, an arrival on a server that has i jobs is allocated to this server with probability $Q_i(t) = (Y_{ii}(t)/2 + \sum_{j=i+1}^{\infty} Y_{ij}(t))/X_i(t)$.

We now look at the evolution of $Y_{ij}(t)$. Let (i, j) be the state of a pair of servers connected by an edge. This state becomes $(i - 1, j)$ when there is a departure on i , which occurs at rate 1 if $i \geq 1$. It becomes $(i + 1, j)$ when there is an arrival on i . This can be caused by two types of events:

- (a) arrival on the edge (i, j) – if each node has k neighbors, an edge (i, j) is chosen at rate $2\rho/k$ and the packet is allocated to the first server with probability $a(i, j) = 1$ if $i < j$, $a(i, i) = 1/2$ and $a(i, j) = 0$ if $i > j$
- (b) arrival on another neighbor of the first server – each other neighbor of i that has state ℓ induces an arrival on i at rate $2\rho a(i, \ell)/k$. Let $Z_{\ell, i, j}(t)$ be the proportion of connected triplets of stations having state (ℓ, i, j) . The arrivals on the first server of a pair (i, j) from one of the $k - 1$ other neighbors occur at rate $2\rho(k - 1)R_{ij}(t)/k$, where $R_{ij}(t) = (Z_{i, i, j}(t)/2 + \sum_{\ell=i+1}^{\infty} Z_{\ell, i, j}(t))/Y_{ij}(t)$.

This shows that, as $X(t)$, the process $Y(t)$ is not a density dependent process because the rates of its transitions involve quantities that depend on triplets. In what follows, we consider a density dependent population process that is an approximation of the original process and has the same transitions but with different rates: We approximate $Z_{\ell, i, j}$ by $Y_{i, \ell}(t)Y_{i, j}(t)/X_i(t)$, which amounts at replacing $R_{ij}(t)$ by $Q_i(t)$. This approximation is called the pair-approximation and leads to the following differential equation for y_{ij} :

$$\begin{aligned} \dot{y}_{ij} = & \left(y_{i+1, j} - y_{i, j} \mathbf{1}_{\{i>0\}} + y_{i, j+1} - y_{i, j} \mathbf{1}_{\{j>0\}} \right) \\ & + \frac{2\rho}{k} (y_{i-1, j} a(i-1, j) + y_{i, j-1} a(j-1, i) - y_{ij}) \\ & + \frac{2\rho(k-1)}{k} (q_{i-1} y_{i-1, j} \mathbf{1}_{\{i>0\}} + q_{j-1} y_{i, j-1} \mathbf{1}_{\{j>0\}} - (q_i + q_j) y_{ij}) \end{aligned} \quad (5.2)$$

The first line of this equation corresponds to the rate of changes of the proportion of pairs (i, j) induced by the departures; the second line is for the arrival on the pair (i, j) and the last line on the arrival on the neighbors or i and j . The first two lines are exact while the last involves the approximation $z_{\ell, i, j} \approx y_{i, \ell} y_{i, j} / x_i$.

5.2.4 Numerical Evaluation

We now compare numerically the steady state distribution (computed by simulation) with the mean field approximation and the fixed-point of the pair-approximation ODE (5.2). Note that we did the comparison for values from $\rho = 0.5$ to $\rho = 0.99$ and only a subset of the results are reported here. All tested values show that the pair-approximation provides an excellent approximation of the shape of the stationary distribution.

The computation of the steady-state distribution is obtained by running a discrete-event simulator that follows strictly the model. In all cases, we simulate a system with $N = 1000$ servers for a total $T = 10^{11}$ events. Comparisons with smaller values of T indicate that $T = 10^{11}$ is enough to reach the steady-state. The fixed point of the pair-approximation equations is computed by integrating numerically the system of differential equations (5.2).

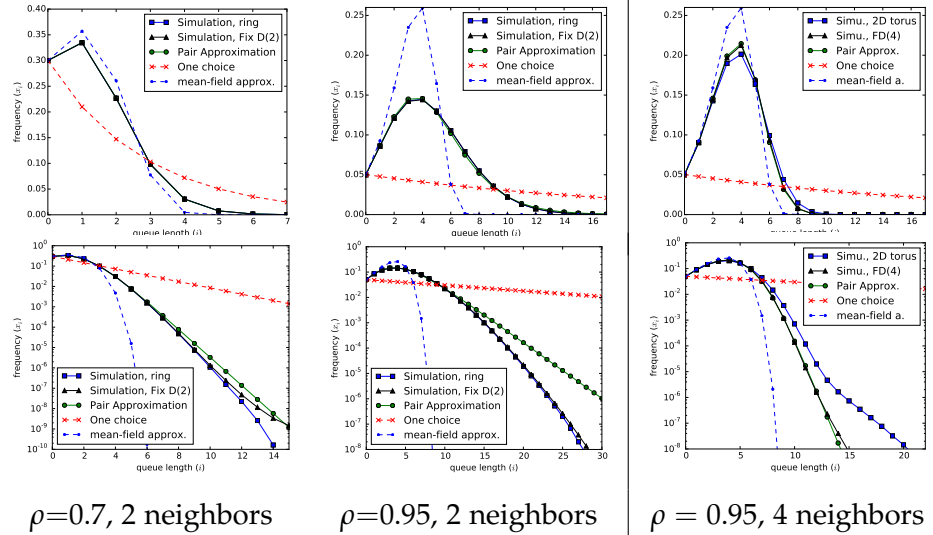


Figure 5.4: Steady-state probability for a server to have i jobs as a function of i . We compare values obtained by simulations and two fluid approximations (mean field and pair-approximation). Second line is in log-scale.

In Figure 5.4, we report the steady-state probability x_i that a given server has i jobs as a function of i . Each plot compares five curves: two are obtained by simulation – (ring/random graph with fixed degree $k = 2$) for the first two plots and (2D torus/random graph with $k = 4$) for the last one –, and two are the fluid approximations of the model (mean field and pair approximation with $k = 2$ or $k = 4$). The last curve corresponds to a model without choices (each server is an independent M/M/1 queue) and is here for comparison. These results show that the pair-approximation predicts very accurately the general shape of the distribution of the simulated model, which are far from both the one-choice and the mean field approximation. The tail of the distribution, however, does not seem to be correct, even if it is much closer for the pair-approximation than for the mean field model.

This figure shows that pair approximation provides a very good estimate of the shape of the distribution but does not predict well the tail. Note that this is similar to what happens for the refined mean field approximation.

5.3 FORECASTING VIA MEAN FIELD APPROXIMATION

In [P14], we study the problem of making forecasts about the future availability of bicycles in stations of a bike-sharing system (BSS). This is relevant in order to make recommendations guaranteeing that the probability that a user will be able to make a journey is sufficiently high. To do this we use probabilistic predictions obtained from a queuing theoretical time-inhomogeneous model of a BSS. The use of a

Recall that the mean field approximation is very accurate for the complete graph but not here because the graph that we consider are very sparse.

model of independent queue is similar to the mean field approximation that we studied in Section 5.1. One of the main contribution of [P14] is to develop a critique of the standard root-mean-square-error (RMSE) – commonly adopted in the bike-sharing research as an index of the prediction accuracy. We believe that RMSE is not appropriate to estimate the quality of forecasts in bike-sharing systems because RMSE is too sensitive to the stochasticity inherent in the real system. Instead we show that using metrics based perfect scoring rules is more appropriate. The model is parametrized and successfully validated using historical data from the Velib BSS of the city of Paris.

5.3.1 Model and Validation

We consider a model in which each station is considered to be independent of its neighbors. We assume that users arrive at a station i (to pick up a bike) according to a Poisson process of time-dependent intensity $\lambda_i(t)$. Similarly, user returns bike at the station i at rate $\mu_i(t)$. This model makes two main approximations. First, this model assumes that arrival of bikes and user can be represented by Poisson processes. Second, this model neglects that neighboring stations have an effect on the bikes returned in station i . The model could be modified to better take this into account but this leads to little improvement in forecasting quality [30].

In all numerical evaluation, we estimate the values of $\lambda_i(t)$ and $\mu_i(t)$ by using historical data from the *Vélib'* system in Paris collected between 1 October 2013 and 31 December 2014 and provided by the operator. We use this information to compute a historical predictor and estimate the (bike) departure and arrival rates, for each minute of a typical working day or weekend day (see the details in the next section). During the day, $\lambda_i(t)$ and $\mu_i(t)$ seem to be of the order of 5 bikes per hour, although there are huge disparities between stations.

The use of Poisson processes to model arrivals in bike-sharing systems has been widely used in the literature (see for instance [36], [P7] and the paper that cite them). Also, it model is partially supported by some statistical tests, see [P14, §3.2].

5.3.2 Deterministic Forecasts are not Sufficient

Most of the papers dealing with prediction for bike-sharing systems, e.g., [34], [41], [51], [97], [94], focus on deterministic forecasts: for a given time horizon h , the challenge is to predict at time $t + h$ a single value for $x_t(t + h)$ that is as close as possible to the “real” value $x(t + h)$. The performance is generally evaluated using the root-mean-square-error (RMSE):

$$\text{RMSE} = \sqrt{\mathbb{E} \left[\|x_t(t + h) - x(t + h)\|^2 \right]}$$

This metric is the most commonly used in prediction challenges [68].

We argue in [P14] that this metric is non well-suited for bike-sharing systems because the system has too much randomness. More specif-

ically, we show that regardless of their quality, the RMSE of such predictors will always be large due to the stochastic nature of bike-sharing systems. In particular, our results suggest that, for a prediction horizon of $h = 1$ hour, there exist no predictor that has a RMSE smaller than 3.5 bikes. In fact, the authors of [97] compare various deterministic forecast that all attain a performance that equals our bounds plus a few percents. Our result suggest that these forecasts are all very close to optimal.

5.3.3 Probabilistic Forecasts

As an alternative to deterministic forecast, we propose to use probabilistic forecasts: for a given time horizon h , the challenge is to predict a probability distribution $P_t(t+h)$ for $X(t+h)$ given the information available at time t . We illustrate the difference between deterministic and probabilistic forecasts in Figure 5.5. All forecasts are issued at 7am. The solid green curve corresponds to the number of bikes in the station up to 7am. Figure 5.5a shows three examples of deterministic forecasts that are classically used in the literature: a first that only used the historical data and ignore the green curve (history), a second that only observe the station state at 7am (persistent forecast) and a third that combined the two notion (best deterministic). Figure 5.5b shows an example of probabilistic forecast where all the potential future distribution of occupancy is predicted at 7am.

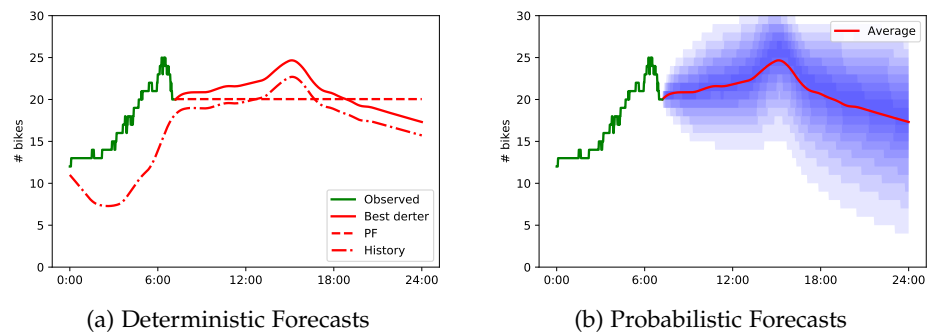


Figure 5.5: Illustration of the two types of forecasts.

One of the main difficulties when dealing with probabilistic forecasts is to evaluate their quality: for each prediction of $P_t(t+h)$, only one realization of $X(t+h)$ will be available. The most natural way to evaluate a forecast is to use proper scoring rules. One of the most well known proper scoring rule is the log-likelihood: the score received is the log of the likelihood of the prediction. A proper scoring rule encourages honesty and will encourage the forecaster to reveal the confidence that she has in her forecast. A very confident forecaster will issue a probability distribution that is very concentrated near a single value x . In this case, the score will be high when the observation

$X \approx x$ but very low when X is far from x . On the other hand, if the forecaster has a low confidence, she will issue a forecast that spans many values. The received score will neither be very high or very low.

There exist many variants of proper scoring rule. We refer to [37] for a complete and clear exposition on the subject. In [P14] we compare different proper scoring rules that are more appropriate to bike sharing systems than the log-likelihood. We also show that the forecast based on our queuing model performs well.

5.4 EXTENSION AND OPEN QUESTIONS

Vehicle sharing systems are now widely developed around the world and similar means of transportation are rapidly growing (see for instance the recent deployment of electric scooters in France). These systems induce complex saturation dynamics (unavailable vehicles or parking space). There is nowadays an active field of research on how to design efficient redistribution or pricing policies to improve the performance of such systems [15], [25], [59], [70], [77]. These papers often focus on a single system (for instance, New York in [25]) and it is not clear whether the good performance of the proposed heuristics comes from a particular feature of this case study or it performs well in general.

This calls for a sound and reproducible methodology to develop and evaluate control heuristics that are widely applicable. Our research agenda aims at obtaining a better understanding of vehicle sharing systems by proposing clean theoretical models, new optimization algorithms and a reproducible simulation environment that can validate the performance of the proposed heuristics. This leads to a number of research challenges.

- **Geometric mean field models** – Mean field approximation provide an accurate description of systems composed of objects in which an individual can interact with many other agents. This approximation also works when the graph of interaction is not complete but just dense as in [69]. When the graph has a small degree, pair approximation seems to provide a refined geometric approximation but for now a justification of this approximation is missing. Such a justification could be done by using tools from probabilistic stochastic automata [64] or the recent advances of [46].
- **Optimization algorithms for vehicle sharing systems** – The performance of vehicle sharing systems rely on good control heuristics (for instance concerning vehicle relocation and/or demand segregation). The development of those classically rely on tools from deterministic operations research problems (e.g., vehicle routing problems) but do not always perform well in

a system as random as a bike-sharing system. To tackle this issue, we plan to use several recent theoretical development in distributed stochastic optimization such as decomposition methods [20], mean field control [31] or refined mean field optimization [P23].

- **Reproducible simulation environment** – We believe that the literature on bike-sharing systems lacks of benchmarks. In order to validate (or invalidate) that our proposed heuristics perform well in real systems, we plan to develop an open source simulator in which it will be easy to program and assess new and existing control strategies. Our aim is to be able to evaluate these strategies on realistic dataset. For that, we will use existing traces of usage of existing systems (e.g. [89]) and combine this with traces analysis and data-mining expertise to create realistic scenarios.

We believe that, given the growing importance of this mode of transportation, these challenges are not only challenging research problems but can also have a societal impact. In addition to the most technical aspects, tackling them will require an integration of tools from data analysis, demand modeling, infrastructure description, and geographic visualization to identify bottlenecks that limit the performance of control strategies of vehicle sharing systems.

CONCLUSION AND PERSPECTIVES

This document summarizes a few of the main contributions of my research. The central theme of this line of research has been to obtain a better understanding of what is mean field approximation. Motivated by applications, a particular emphasis has been to increase the range of application of mean field approximation methods, by showing for instance that:

- There exists a generic ways to deal with non-continuous drifts by using a differential inclusion approach (Chapter 2);
- Mean field approximation can be refined for small system size by using expansion (Chapter 3);
- Mean field (and refined mean field) can deal with heterogeneous models (Chapter 4).
- Pair approximation, a variant of mean field approximation, can be used to study the impact of spatial interactions (Chapter 5).

The strategy to obtain these results varied in time and reflects my understanding of the subject: The first results that I developed were obtained by using sample-path arguments and stochastic approximation methods; The more recent papers are based on Stein's method and the convergence of generators. The later approach provides much tighter bounds that, I believe, are a better justification of the success of mean field approximation.

In order to obtain a concise and uniform manuscript, this document does not covers all my contributions to the development of mean field approximation. The following chapter could have been integrated:

- Mean field games: In [P5], [P4], we develop a generic framework for discrete-space mean field games. We show in particular that mean field games are not always a good approximation of symmetric games with a large number of players: mean field equilibria describe a subset of the original stochastic games.
- Hitting times. In [P21], we show how to use mean field approximation to compute hitting time of a stochastic process. We show that the hitting time of the stochastic system with N objects is close to time for its mean field approximation to be at distance $1/N$ of its fixed point.
- Centralized optimization. In [P8], [P13], [P9], we show that the control of a system with N objects converges to the solution of

a mean field optimal control. More importantly, we show the mean field optimal policy is also asymptotically optimal for the system of size N . This simplifies the computation of an optimal policy.

Several open research questions have already been listed at the end of each chapter and we will not recall them here. In fact, most of these questions are related to a bigger research agenda which is to develop efficient heuristics for the control of stochastic distributed agents. Restless bandit allocation is one particular example where the control that can be sent to each arm is restricted to an on/off signal. We believe that our results on the *refined mean field* will allow the development of control heuristics that are asymptotically optimal as the number of arms goes to infinity and that also have a better performance than existing heuristics for a moderate number of arms. A typical application of this approach would be in the context of smart grids, to develop control policies for distributed electric appliances. Beyond this application, the notion of refined mean field games will help to design efficient allocation, for example in the context of wireless networks where mean field games are already used.

One of the classical approaches to build such heuristics is to use Lagrangian relaxations of the problems to decompose a large-scale optimization problem into smaller independent sub-problems. Depending on the relaxation used, one obtains radically different methods: Whittle index for restless bandits [91], [92], stochastic decomposition methods [20] and SDDP [75]; progressive hedging [73],... Similar to mean field approximation, these methods have been shown to work well in practice and in some cases can be shown to be asymptotically optimal [91], [87].

We believe that, by using our expertise on mean field approximation and control, it will be possible to make connection between bandit optimal control, stochastic decomposition methods and mean field approximation.

REFERENCES

PUBLICATIONS FROM THE AUTHOR

- [P1] L. Bortolussi and N. Gast. “Mean Field Approximation of Uncertain Stochastic Models”. In: *46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2016)*. 2016 (cit. on pp. 2, 5).
- [P2] L. Bortolussi and N. Gast. “Mean-field limits beyond ordinary differential equations”. In: *Formal Methods for the Quantitative Evaluation of Collective Adaptive Systems*. Springer, 2016, pp. 61–82 (cit. on pp. 2, 5, 9, 13).
- [P3] G. Casale and N. Gast. “Asymptotic analysis methods for list-based caches with non-uniform access costs”. 2019 (cit. on pp. 33, 36).
- [P4] J. Doncel, N. Gast, and B. Gaujal. “Are mean-field games the limits of finite stochastic games?” In: *The 18th Workshop on Mathematical Performance Modeling and Analysis (MAMA)*. 2016 (cit. on p. 57).
- [P5] J. Doncel, N. Gast, and B. Gaujal. “Discrete mean field games: Existence of equilibria and convergence”. In: *Journal of Dynamics & Games* (2019), pp. 269–316 (cit. on p. 57).
- [P6] C. Fricker, N. Gast, and A. Mohamed. “Mean field analysis for inhomogeneous bike sharing systems”. In: *AofA 2012, International Meeting on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms* (2012) (cit. on pp. 3, 46, 48).
- [P7] C. Fricker and N. Gast. “Incentives and redistribution in homogeneous bike-sharing systems with stations of finite capacity”. In: *Euro journal on transportation and logistics* 5.3 (2016), pp. 261–291 (cit. on pp. 3, 12, 45–48, 53).
- [P8] N. Gast and B. Gaujal. “A mean field approach for optimization in discrete time”. In: *Discrete Event Dynamic Systems* (2011). ISSN: 0924-6703 (cit. on pp. 13, 57).
- [P9] N. Gast and B. Gaujal. “A Mean Field Approach for Optimization in Particle Systems and Applications”. In: *Fourth International Conference on Performance Evaluation Methodologies and Tools, ValueTools* (2009). Best Student Paper Award (cit. on p. 57).
- [P10] N. Gast and B. Gaujal. “A Mean Field Model of Work Stealing in Large-Scale Systems”. In: *ACM SIGMETRICS* (2010) (cit. on pp. 9, 25).
- [P11] N. Gast and B. Gaujal. “Markov chains with discontinuous drifts have differential inclusion limits”. In: *Performance Evaluation* (2012) (cit. on pp. 2, 5, 11, 13, 17, 19, 48).
- [P12] N. Gast and B. Gaujal. “Mean field limit of non-smooth systems and differential inclusions”. In: *ACM SIGMETRICS Performance Evaluation Review* 38.2 (2010), pp. 30–32 (cit. on pp. 2, 5).

- [P13] N. Gast, B. Gaujal, and J. Le Boudec. “Mean field for Markov decision processes: from discrete to continuous optimization”. In: *IEEE Transactions on Automatic Control* 57.9 (2012), pp. 2266–2280 (cit. on p. 57).
- [P14] N. Gast, G. Massonnet, D. Reijsbergen, and M. Tribastone. “Probabilistic forecasts of bike-sharing systems for journey planning”. In: *Proceeding of the 24th ACM International Conference on Information and Knowledge Management (CIKM’15)*. ACM, 2015 (cit. on pp. 3, 45, 52, 53, 55).
- [P15] N. Gast. “Construction of Lyapunov functions via relative entropy with application to caching”. In: *The 18th Workshop on MATHematical performance Modeling and Analysis (MAMA)*. 2016 (cit. on pp. 3, 9, 12, 33, 37).
- [P16] N. Gast. “Expected Values Estimated via Mean-Field Approximation Are $1/N$ -Accurate”. In: *Proc. ACM Meas. Anal. Comput. Syst.* 1.1 (June 2017), 17:1–17:26. ISSN: 2476-1249. DOI: [10.1145/3084454](https://doi.org/10.1145/3084454). URL: <http://doi.acm.org/10.1145/3084454> (cit. on pp. 3, 15, 17, 19).
- [P17] N. Gast. “Expected Values Estimated via Mean-Field Approximation Are $1/N$ -Accurate: Extended Abstract”. In: *Proceedings of the 2017 ACM SIGMETRICS / International Conference on Measurement and Modeling of Computer Systems*. SIGMETRICS ’17 Abstracts. Urbana-Champaign, Illinois, USA: ACM, 2017, pp. 50–50. ISBN: 978-1-4503-5032-7. DOI: [10.1145/3078505.3078523](https://doi.org/10.1145/3078505.3078523). URL: <http://doi.acm.org/10.1145/3078505.3078523> (cit. on p. 3).
- [P18] N. Gast. *Refined Mean Field Tool*. https://github.com/ngast/rmf_tool. 2018 (cit. on pp. 23, 31, 40).
- [P19] N. Gast. “The Power of Two Choices on Graphs: the Pair-Approximation is Accurate”. In: *ACM SIGMETRICS Performance Evaluation Review* (2015) (cit. on pp. 3, 45, 49).
- [P20] N. Gast, L. Bortolussi, and M. Tribastone. “Size Expansions of Mean Field Approximation: Transient and Steady-State Analysis.” In: *Performance Evaluation* (2018) (cit. on pp. 3, 15, 20–23, 28, 29).
- [P21] N. Gast and B. Gaujal. “Computing absorbing times via fluid approximations”. In: *Advances in Applied Probability* 49.3 (2017), pp. 768–790. DOI: [10.1017/apr.2017.21](https://doi.org/10.1017/apr.2017.21) (cit. on p. 57).
- [P22] N. Gast, D. Latella, and M. Massink. “A refined mean field approximation of synchronous discrete-time population models”. In: *Performance Evaluation* 126 (2018), pp. 1–21. ISSN: 0166-5316. DOI: <https://doi.org/10.1016/j.peva.2018.05.002> (cit. on pp. 3, 15, 31).
- [P23] N. Gast and B. Van Houdt. “A Refined Mean Field Approximation”. In: *Proc. ACM Meas. Anal. Comput. Syst.* 1.2 (Dec. 2017), 33:1–33:28. ISSN: 2476-1249. DOI: [10.1145/3154491](https://doi.org/10.1145/3154491). URL: <http://doi.acm.org/10.1145/3154491> (cit. on pp. 3, 15, 20, 24, 28, 41, 56).

- [P24] N. Gast and B. Van Houdt. “A Refined Mean Field Approximation”. In: *Abstracts of the 2018 ACM International Conference on Measurement and Modeling of Computer Systems*. SIGMETRICS '18. Irvine, CA, USA: ACM, 2018, pp. 113–113. ISBN: 978-1-4503-5846-0. DOI: [10.1145/3219617.3219663](https://doi.org/10.1145/3219617.3219663). URL: <http://doi.acm.org/10.1145/3219617.3219663> (cit. on p. 3).
- [P25] N. Gast and B. Van Houdt. “Asymptotically Exact TTL-Approximations of the Cache Replacement Algorithms LRU(m) and h-LRU”. In: *28th International Teletraffic Congress (ITC 28)*. 2016 (cit. on pp. 3, 5, 7, 33, 34, 38, 42).
- [P26] N. Gast and B. Van Houdt. “Transient and Steady-state Regime of a Family of List-based Cache Replacement Algorithms”. In: *Proceedings of ACM SIGMETRICS*. ACM, 2015 (cit. on pp. 3, 5, 7, 33–38).
- [P27] N. Gast and B. Van Houdt. “Transient and steady-state regime of a family of list-based cache replacement algorithms”. In: *Queueing Systems* 83.3-4 (2016), pp. 293–328 (cit. on pp. 3, 7, 33).
- [P28] N. Gast and B. Van Houdt. “TTL approximations of the cache replacement algorithms LRU(m) and h-LRU”. In: *Performance Evaluation* 117 (2017), pp. 33–57. ISSN: 0166-5316. DOI: <https://doi.org/10.1016/j.peva.2017.09.002>. URL: <http://www.sciencedirect.com/science/article/pii/S0166531617300688> (cit. on pp. 3, 33, 42).

REFERENCES

- [1] R. Aghajani, K. Ramanan, et al. “The hydrodynamic limit of a randomized load balancing network”. In: *The Annals of Applied Probability* 29.4 (2019), pp. 2114–2174 (cit. on p. 32).
- [2] A. Ale, P. Kirk, and M. P. Stumpf. “A general moment expansion method for stochastic kinetic models”. In: *The Journal of chemical physics* 138.17 (2013), p. 174101 (cit. on p. 24).
- [3] O. I. Aven, E. G. Coffman Jr., and Y. A. Kogan. *Stochastic Analysis of Computer Storage*. Norwell, MA, USA: Kluwer Academic Publishers, 1987. ISBN: 9027725152 (cit. on pp. 5, 34–36).
- [4] F. Baccelli and T. Taillefumier. “Replica-mean-field limits for intensity-based neural networks”. In: *arXiv preprint arXiv:1902.03504* (2019) (cit. on p. 2).
- [5] M. Benaim. “Dynamics of stochastic approximation algorithms”. In: *Seminaire de probabilites XXXIII*. Springer, 1999, pp. 1–68 (cit. on p. 11).
- [6] M. Benaim. “Recursive algorithms, urn processes and chaining number of chain recurrent sets”. In: *Ergodic Theory and Dynamical Systems* 18.1 (1998), pp. 53–87 (cit. on p. 11).
- [7] M. Benaim and J.-Y. L. Boudec. “On mean field convergence and stationary regime”. In: *arXiv preprint arXiv:1111.5710* (2011) (cit. on p. 12).
- [8] M. Benaim and J.-Y. Le Boudec. “A class of mean field interaction models for computer and communication systems”. In: *Performance Evaluation* 65.11 (2008), pp. 823–838 (cit. on pp. 2, 10–13, 15, 18, 31).

- [9] O. Bencheikh and B. Jourdain. “Bias behaviour and antithetic sampling in mean-field particle approximations of SDEs nonlinear in the sense of McKean”. In: *ESAIM: Proceedings and Surveys* 65 (2019), pp. 219–235 (cit. on p. 16).
- [10] J. van den Berg and A. Gandolfi. “LRU Is Better than FIFO under the Independent Reference Model”. English. In: *Journal of Applied Probability* 29.1 (1992), pp. 239–243. ISSN: 00219002. URL: <http://www.jstor.org/stable/3214811> (cit. on p. 35).
- [11] M. van der Boor, S. C. Borst, J. S. van Leeuwen, and D. Mukherjee. “Scalable load balancing in networked systems: A survey of recent advances”. In: *arXiv preprint arXiv:1806.05444* (2018) (cit. on p. 49).
- [12] V. S. Borkar. *Stochastic approximation: a dynamical systems viewpoint*. Vol. 48. Springer, 2009 (cit. on p. 11).
- [13] L. Bortolussi and R. A. Hayden. “Bounds on the deviation of discrete-time Markov chains from their mean-field model”. In: *Performance Evaluation* 70.10 (2013), pp. 736–749 (cit. on p. 10).
- [14] J.-Y. L. Boudec. “The stationary behaviour of fluid limits of reversible processes is concentrated on stationary points”. In: *arXiv preprint arXiv:1009.5021* (2010) (cit. on p. 12).
- [15] B. Boyacı, K. G. Zografos, and N. Geroliminis. “An optimization framework for the development of efficient one-way car-sharing systems”. In: *European Journal of Operational Research* 240.3 (2015), pp. 718–733 (cit. on p. 55).
- [16] M. Bramson, Y. Lu, and B. Prabhakar. “Randomized load balancing with general service time distributions”. In: *ACM SIGMETRICS performance evaluation review*. Vol. 38. 1. ACM, 2010, pp. 275–286 (cit. on p. 32).
- [17] A. Braverman, J. Dai, et al. “Stein’s method for steady-state diffusion approximations of M/Ph/n+ M systems”. In: *The Annals of Applied Probability* 27.1 (2017), pp. 550–581 (cit. on p. 18).
- [18] A. Braverman, J. Dai, and J. Feng. “Stein’s method for steady-state diffusion approximations: an introduction through the Erlang-A and Erlang-C models”. In: *Stochastic Systems* 6.2 (2017), pp. 301–366 (cit. on p. 18).
- [19] A. S. Budhiraja, P. Dupuis, M. Fischer, and K. Ramanan. “Limits of relative entropies associated with weakly interacting particle systems”. In: *Electronic Journal of Probability* 20 (2015) (cit. on p. 12).
- [20] P. Carpentier, J.-P. Chancelier, M. D. Lara, and T. Rigaut. “Time Blocks Decomposition of Multistage Stochastic Optimization Problems”. In: (Apr. 2018) (cit. on pp. 56, 58).
- [21] G. Casale. “Analyzing replacement policies in list-based caches with non-uniform access costs”. In: *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 432–440 (cit. on p. 42).
- [22] A. Chaintreau, J.-Y. Le Boudec, and N. Ristanovic. “The Age of Gossip: Spatial Mean Field Regime”. In: *SIGMETRICS Perform. Eval. Rev.* 37.1 (June 2009), pp. 109–120. ISSN: 0163-5999. DOI: [10.1145/2492101.1555363](https://doi.org/10.1145/2492101.1555363). URL: <http://doi.acm.org/10.1145/2492101.1555363> (cit. on p. 9).

- [23] H. Che, Y. Tung, and Z. Wang. "Hierarchical Web caching systems: modeling, design and experimental results". In: *IEEE J.Sel. A. Commun.* 20.7 (2002), pp. 1305–1314 (cit. on p. 5).
- [24] J.-w. Cho, J.-Y. Le Boudec, and Y. Jiang. "On the validity of the fixed point equation and decoupling assumption for analyzing the 802.11 mac protocol". In: *ACM SIGMETRICS Performance Evaluation Review* 38.2 (2010), pp. 36–38 (cit. on pp. 9, 12).
- [25] H. Chung, D. Freund, and D. B. Shmoys. "Bike Angels: An Analysis of Citi Bike's Incentive Program". In: *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. ACM. 2018, p. 5 (cit. on pp. 47, 55).
- [26] A. Dan and D. Towsley. *An approximate analysis of the LRU and FIFO buffer replacement schemes*. Vol. 18. 1. ACM, 1990 (cit. on p. 7).
- [27] P. Dupuis and M. Fischer. "On the construction of lyapunov functions for nonlinear markov processes via relative entropy". In: *submitted for publication* (2011) (cit. on p. 12).
- [28] R. Fagin. "Asymptotic miss ratios over independent references". In: *Journal of Computer and System Sciences* 14.2 (1977), pp. 222–250. ISSN: 0022-0000. DOI: [http://dx.doi.org/10.1016/S0022-0000\(77\)80014-7](http://dx.doi.org/10.1016/S0022-0000(77)80014-7). URL: <http://www.sciencedirect.com/science/article/pii/S0022000077800147> (cit. on p. 5).
- [29] R. Fagin and T. G. Price. "Efficient Calculation of Expected Miss Ratios in the Independent Reference Model". In: *SIAM J. Comput.* 7 (1978), pp. 288–296 (cit. on p. 6).
- [30] C. Feng, J. Hillston, and D. Reijnders. "Moment-based probabilistic prediction of bike availability for bike-sharing systems". In: *International Conference on Quantitative Evaluation of Systems*. Springer. 2016, pp. 139–155 (cit. on p. 53).
- [31] M. Fornasier and F. Solombrino. "Mean-field optimal control". In: *ESAIM: Control, Optimisation and Calculus of Variations* 20.4 (2014), pp. 1123–1152 (cit. on p. 56).
- [32] C. M. Fortuin, P. W. Kasteleyn, and J. Ginibre. "Correlation inequalities on some partially ordered sets". In: *Communications in Mathematical Physics* 22.2 (1971), pp. 89–103. URL: <http://projecteuclid.org/euclid.cmp/1103857443> (cit. on p. 35).
- [33] C. Fricker, P. Robert, and J. Roberts. "A versatile and accurate approximation for LRU cache performance". In: *2012 24th International Teletraffic Congress (ITC 24)*. IEEE. 2012, pp. 1–8 (cit. on pp. 38, 39).
- [34] J. Froehlich, J. Neumann, and N. Oliver. "Sensing and Predicting the Pulse of the City through Shared Bicycling". In: *IJCAI*. Vol. 9. 2009, pp. 1420–1426 (cit. on p. 53).
- [35] E. Gelenbe. "A Unified Approach to the Evaluation of a Class of Replacement Algorithms". In: *IEEE Trans. Comput.* 22.6 (June 1973), pp. 611–618. ISSN: 0018-9340. DOI: [10.1109/TC.1973.5009115](http://dx.doi.org/10.1109/TC.1973.5009115). URL: <http://dx.doi.org/10.1109/TC.1973.5009115> (cit. on pp. 5, 6, 34).
- [36] D. George and C. Xia. "Fleet-sizing and service availability for a vehicle rental system via closed queueing networks". In: *European Journal of Operational Research* 211.1 (2011), pp. 198–207 (cit. on p. 53).

- [37] T. Gneiting and A. E. Raftery. “Strictly proper scoring rules, prediction, and estimation”. In: *Journal of the American Statistical Association* 102.477 (2007), pp. 359–378 (cit. on p. 55).
- [38] J. Goutsias and G. Jenkinson. “Markovian dynamics on complex reaction networks”. In: *Physics Reports* 529.2 (2013), pp. 199–264 (cit. on p. 24).
- [39] R. Grima. “A study of the accuracy of moment-closure approximations for stochastic chemical kinetics”. In: *The Journal of chemical physics* 136.15 (2012), 04B616 (cit. on pp. 24, 25).
- [40] R. Grima. “An effective rate equation approach to reaction kinetics in small volumes: Theory and application to biochemical reactions in nonequilibrium steady-state conditions”. In: *The Journal of chemical physics* 133.3 (2010), 07B604 (cit. on p. 24).
- [41] M. C. Guenther and J. T. Bradley. “Journey data based arrival forecasting for bicycle hire schemes”. In: *Analytical and Stochastic Modeling Techniques and Applications*. Springer, 2013, pp. 214–231 (cit. on p. 53).
- [42] M. C. Guenther, A. Stefanek, and J. T. Bradley. “Moment closures for performance models with highly non-linear rates”. In: *European Workshop on Performance Engineering*. Springer, 2012, pp. 32–47 (cit. on p. 24).
- [43] I. Gurvich. “Diffusion models and steady-State approximations for exponentially ergodic Markovian queues”. In: *Annals of Applied Probability* 24.6 (2014), pp. 2527–2559 (cit. on p. 18).
- [44] J. H. Hester and D. S. Hirschberg. “Self-organizing Linear Search”. In: *ACM Comput. Surv.* 17.3 (Sept. 1985), pp. 295–311. ISSN: 0360-0300. DOI: [10.1145/5505.5507](https://doi.org/10.1145/5505.5507). URL: <http://doi.acm.org/10.1145/5505.5507> (cit. on p. 34).
- [45] R. Hirade and T. Osogami. “Analysis of Page Replacement Policies in the Fluid Limit”. In: *Oper. Res.* 58 (July 2010), pp. 971–984. ISSN: 0030-364X. DOI: [10.1287/opre.1090.0761](https://doi.org/10.1287/opre.1090.0761). URL: <http://dx.doi.org/10.1287/opre.1090.0761> (cit. on p. 7).
- [46] L. Hodgkinson, R. McVinish, and P. K. Pollett. “Normal approximations for discrete-time occupancy processes”. In: *arXiv preprint arXiv:1801.00542* (2018) (cit. on pp. 14, 55).
- [47] M. Huang, R. P. Malhamé, P. E. Caines, et al. “Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle”. In: *Communications in Information & Systems* 6.3 (2006), pp. 221–252 (cit. on p. 7).
- [48] P. Jelenkovic. “Asymptotic approximation of the move-to-front search cost distribution and least-recently used caching fault probabilities”. In: *Ann. Appl. Probab.* 9.2 (May 1999), pp. 430–464. DOI: [10.1214/aop/1029962750](https://doi.org/10.1214/aop/1029962750). URL: <http://dx.doi.org/10.1214/aop/1029962750> (cit. on p. 7).
- [49] B. Jiang, P. Nain, and D. Towsley. “On the Convergence of the TTL Approximation for an LRU Cache under Independent Stationary Request Processes”. In: *ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS)* 3.4 (2018), p. 20 (cit. on pp. 38, 39).

- [50] E. Jones, T. Oliphant, P. Peterson, et al. *SciPy: Open source scientific tools for Python*. [Online; accessed 2017-06-15]. 2001. URL: <http://www.scipy.org/> (cit. on p. 23).
- [51] A. Kaltenbrunner, R. Meza, J. Grivolla, J. Codina, and R. Banchs. “Urban cycles and mobility patterns: Exploring and predicting trends in a bicycle-based public transport system”. In: *Pervasive and Mobile Computing* 6.4 (2010), pp. 455–466 (cit. on p. 53).
- [52] K. Kenthapadi and R. Panigrahy. “Balanced allocation on graphs”. In: *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. SIAM. 2006, pp. 434–443 (cit. on p. 49).
- [53] W. F. King III. “Analysis of Demand Paging Algorithms”. In: *IFIP Congress (1)*. 1971, pp. 485–490 (cit. on pp. 6, 34).
- [54] V. N. Kolokoltsov, J. Li, and W. Yang. “Mean field games and non-linear Markov processes”. In: *arXiv preprint arXiv:1112.3744* (2011) (cit. on pp. 10, 18, 19).
- [55] C. Kuehn. “Moment closure—a brief review”. In: *Control of self-organizing nonlinear systems*. Springer, 2016, pp. 253–271 (cit. on p. 24).
- [56] T. G. Kurtz. “Solutions of Ordinary Differential Equations as Limits of Pure Jump Markov Processes”. In: *Journal of Applied Probability* 7 (1970), pp. 49–58 (cit. on pp. 8, 10, 13, 15).
- [57] T. G. Kurtz. “Strong approximation theorems for density dependent Markov chains”. In: *Stochastic Processes and Their Applications* 6.3 (1978), pp. 223–240 (cit. on pp. 2, 8).
- [58] H. Kushner and G. G. Yin. *Stochastic approximation and recursive algorithms and applications*. Vol. 35. Springer Science & Business Media, 2003 (cit. on p. 11).
- [59] G. Laporte, F. Meunier, and R. W. Calvo. “Shared mobility systems”. In: *4or* 13.4 (2015), pp. 341–360 (cit. on p. 55).
- [60] J.-M. Lasry and P.-L. Lions. “Mean field games”. In: *Japanese journal of mathematics* 2.1 (2007), pp. 229–260 (cit. on p. 2).
- [61] D. Latella, M. Loreti, and M. Massink. “On-the-fly PCTL fast mean-field approximated model-checking for self-organising coordination”. In: *Science of Computer Programming* 110 (2015), pp. 23–50 (cit. on p. 14).
- [62] J.-Y. Le Boudec, D. McDonald, and J. Munding. “A generic mean field convergence result for systems of interacting objects”. In: *Fourth international conference on the quantitative evaluation of systems (QEST 2007)*. IEEE. 2007, pp. 3–18 (cit. on pp. 2, 13).
- [63] Y. Lu, Q. Xie, G. Kliot, A. Geller, J. R. Larus, and A. Greenberg. “Join-Idle-Queue: A novel load balancing algorithm for dynamically scalable web services”. In: *Performance Evaluation* 68.11 (2011), pp. 1056–1071 (cit. on pp. 9, 25).
- [64] I. Marcovici. “Automates cellulaires probabilistes et mesures spécifiques sur des espaces symboliques”. PhD thesis. Université Paris-Diderot-Paris VII, 2013 (cit. on p. 55).

- [65] S. Méléard. “Asymptotic behaviour of some interacting particle systems; McKean-Vlasov and Boltzmann models”. In: *Probabilistic models for nonlinear partial differential equations*. Springer, 1996, pp. 42–95 (cit. on p. 7).
- [66] W. Minnebo and B. Van Houdt. “A fair comparison of pull and push strategies in large distributed networks”. In: *IEEE/ACM Transactions on Networking (TON)* 22.3 (2014), pp. 996–1006 (cit. on pp. 9, 25–27).
- [67] M. D. Mitzenmacher. “The Power of Two Random Choices in Randomized Load Balancing”. PhD thesis. PhD thesis, Graduate Division of the University of California at Berkeley, 1996 (cit. on pp. 2, 9, 12, 25–27, 49).
- [68] *MoReBikeS: Model Reuse with Bike rental Station data*. <http://reframed2k.org/Challenge> and <http://www.ecmlpkdd2015.org/discovery-challenges>. Accessed 17/04/2015 (cit. on p. 53).
- [69] D. Mukherjee, S. C. Borst, J. S. Van Leeuwen, and P. A. Whiting. “Universality of power-of-d load balancing in many-server systems”. In: *Stochastic Systems* 8.4 (2018), pp. 265–292 (cit. on pp. 49, 55).
- [70] E. O’Mahony and D. B. Shmoys. “Data analysis and optimization for (citi) bike sharing”. In: *Twenty-ninth AAAI conference on artificial intelligence*. 2015 (cit. on p. 55).
- [71] M. A. M. de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. “Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making”. In: *Swarm Intelligence* 5.3-4 (2011), pp. 305–327 (cit. on p. 14).
- [72] H. Ohtsuki, C. Hauert, E. Lieberman, and M. A. Nowak. “A simple rule for the evolution of cooperation on graphs and social networks”. In: *Nature* 441.7092 (2006), pp. 502–505 (cit. on p. 49).
- [73] R. T. Rockafellar and R. J.-B. Wets. “Scenarios and policy aggregation in optimization under uncertainty”. In: *Mathematics of operations research* 16.1 (1991), pp. 119–147 (cit. on p. 58).
- [74] D. Schnoerr, G. Sanguinetti, and R. Grima. “Comparison of different moment-closure approximations for stochastic chemical kinetics”. In: *The Journal of Chemical Physics* 143.18 (2015), 11B610_1 (cit. on p. 24).
- [75] A. Shapiro. “Analysis of stochastic dual dynamic programming method”. In: *European Journal of Operational Research* 209.1 (2011), pp. 63–72 (cit. on p. 58).
- [76] A. Singh and J. P. Hespanha. “Approximate moment dynamics for chemically reacting systems”. In: *IEEE Transactions on Automatic Control* 56.2 (2011), pp. 414–418 (cit. on p. 24).
- [77] D. Singhvi, S. Singhvi, P. I. Frazier, S. G. Henderson, E. O’Mahony, D. B. Shmoys, and D. B. Woodard. “Predicting bike usage for new york city’s bike sharing system”. In: *Workshops at the Twenty-Ninth AAAI Conference on Artificial Intelligence*. 2015 (cit. on p. 55).
- [78] P. Smadbeck and Y. N. Kaznessis. “A closure scheme for chemical master equations”. In: *Proceedings of the National Academy of Sciences* 110.35 (2013), pp. 14261–14265 (cit. on p. 24).

- [79] D. Starobinski and D. Tse. “Probabilistic Methods for Web Caching”. In: *Perform. Eval.* 46.2-3 (Oct. 2001), pp. 125–137. ISSN: 0166-5316. DOI: [10.1016/S0166-5316\(01\)00045-1](https://doi.org/10.1016/S0166-5316(01)00045-1). URL: [http://dx.doi.org/10.1016/S0166-5316\(01\)00045-1](http://dx.doi.org/10.1016/S0166-5316(01)00045-1) (cit. on p. 34).
- [80] C. Stein. “Approximate computation of expectations”. In: *Lecture Notes-Monograph Series* 7 (1986), pp. i–164 (cit. on pp. 10, 18).
- [81] P. Tinnakornsriruphap and A. M. Makowski. “Limit Behavior of ECN/RED Gateways Under a Large Number of TCP Flows”. In: *Proceedings IEEE INFOCOM 2003, The 22nd Annual Joint Conference of the IEEE Computer and Communications Societies, San Francisco, CA, USA, March 30 - April 3, 2003*. IEEE, 2003, pp. 873–883. DOI: [10.1109/INFCOM.2003.1208925](https://doi.org/10.1109/INFCOM.2003.1208925). URL: <https://doi.org/10.1109/INFCOM.2003.1208925> (cit. on p. 13).
- [82] J. N. Tsitsiklis and K. Xu. “On the power of (even a little) centralization in distributed processing”. In: *ACM SIGMETRICS Performance Evaluation Review* 39.1 (2011), pp. 121–132 (cit. on pp. 9, 12, 25).
- [83] N. Tsukada, R. Hirade, and N. Miyoshi. “Fluid Limit Analysis of FIFO and RR Caching for Independent Reference Models”. In: *Perform. Eval.* 69.9 (Sept. 2012), pp. 403–412. ISSN: 0166-5316. DOI: [10.1016/j.peva.2012.05.008](https://doi.org/10.1016/j.peva.2012.05.008). URL: <http://dx.doi.org/10.1016/j.peva.2012.05.008> (cit. on pp. 7, 12).
- [84] B. Van Houdt. “A mean field model for a class of garbage collection algorithms in flash-based solid state drives”. In: *ACM SIGMETRICS Performance Evaluation Review*. Vol. 41. 1. ACM, 2013, pp. 191–202 (cit. on p. 9).
- [85] B. Van Houdt. “Global attraction of ODE-based mean field models with hyperexponential job sizes”. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 3.2 (2019), p. 23 (cit. on p. 12).
- [86] T. Vasantam, A. Mukhopadhyay, and R. R. Mazumdar. “Insensitivity of the mean-field Limit of Loss Systems Under Power-of-d Routing”. In: *arXiv preprint arXiv:1708.09328* (2017) (cit. on p. 48).
- [87] I. M. Verloop. “Asymptotically optimal priority policies for indexable and nonindexable restless bandits”. In: *Ann. Appl. Probab.* 26.4 (Aug. 2016), pp. 1947–1995. DOI: [10.1214/15-AAP1137](https://doi.org/10.1214/15-AAP1137). URL: <https://doi.org/10.1214/15-AAP1137> (cit. on p. 58).
- [88] I. M. Verloop et al. “Asymptotically optimal priority policies for indexable and nonindexable restless bandits”. In: *The Annals of Applied Probability* 26.4 (2016), pp. 1947–1995 (cit. on p. 17).
- [89] *VLS & Stats*. <http://vlsstats.ifsttar.fr/> (cit. on p. 56).
- [90] N. D. Vvedenskaya, R. L. Dobrushin, and F. I. Karpelevich. “Queueing system with selection of the shortest of two queues: An asymptotic approach”. In: *Problemy Peredachi Informatsii* 32.1 (1996), pp. 20–34 (cit. on pp. 9, 25, 26, 49).
- [91] R. R. Weber and G. Weiss. “On an index policy for restless bandits”. In: *Journal of Applied Probability* 27.3 (1990), pp. 637–648 (cit. on p. 58).
- [92] P. Whittle. “Restless bandits: Activity allocation in a changing world”. In: *Journal of applied probability* 25.A (1988), pp. 287–298 (cit. on p. 58).

- [93] Y. Yang, R. Luo, M. Li, M. Zhou, W. Zhang, and J. Wang. “Mean Field Multi-Agent Reinforcement Learning”. In: *Proceedings of the 35th International Conference on Machine Learning*. Ed. by J. Dy and A. Krause. Vol. 80. Proceedings of Machine Learning Research. Stockholm: PMLR, July 2018, pp. 5571–5580. URL: <http://proceedings.mlr.press/v80/yang18d.html> (cit. on p. 2).
- [94] Z. Yang, J. Hu, Y. Shu, P. Cheng, J. Chen, and T. Moscibroda. “Mobility modeling and prediction in bike-sharing systems”. In: *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM. 2016, pp. 165–178 (cit. on p. 53).
- [95] L. Ying. “On the Approximation Error of Mean-Field Models”. In: *Proceedings of the 2016 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Science*. ACM. 2016, pp. 285–297 (cit. on pp. 10, 15, 16, 18, 19).
- [96] L. Ying. “Stein’s Method for Mean Field Approximations in Light and Heavy Traffic Regimes”. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 1.1 (2017), p. 12 (cit. on p. 18).
- [97] J. W. Yoon, F. Pinelli, and F. Calabrese. “Cityride: a predictive bike sharing journey advisor”. In: *Mobile Data Management (MDM), 2012 IEEE 13th International Conference on*. IEEE. 2012, pp. 306–311 (cit. on pp. 53, 54).