École Doctorale de l'École Polytechnique

INRIA PARIS-ROCQUENCOURT

Thèse présentée en vue de l'obtention du grade de Docteur de l'École Polytechnique Spécialité Mathématiques Appliquées

Études de modèles probabilistes de réseaux pair-à-pair et de réseaux avec mobilité

Studies of probabilistic models for peer-to-peer and mobile networks

Florian SIMATOS

First version, July 20, 2009.

Remerciements

Contents

Résumé	vii
Summary	ix
Introduction	
Foreword	1
Modeling of Communication Networks	3
Wireless and Peer-to-Peer Networks	9
Mathematical Framework	18
Presentation of Chapters	28
References of the Introduction	38
Chapter I. Spatial Homogenization in a Stochastic Network with Mobility	
1 Introduction	43
2 Framework and Notations	46
3 Martingale	49
4. Two Key Representations	52
5. The Space Renormalized Process	54
6. The Supercritical Regime	58
7. Stability of the Subcritical Regime	66
Appendix A. Martingale Construction	71
References of Chapter I	81
Chapter II Stability Properties of Linear File-Sharing Networks	
1 Introduction	83
2 Analysis of the Single-Chunk Network	85
3 Vule Processes with Deletions	91
4 Analysis of the Multi-Chunk Network	97
Appendix A Proof of Proposition 3.3	104
References of Chapter II	107
	101
Chapter III. A Queueing System for Modeling a File-Sharing Principle	
1. Introduction	109
2. Model Description	110

3. Bins and Balls Problem
4. Deterministic Problem
5. Random Problem
6. Discussion
References of Chapter III
Chapter IV. Occupancy Schemes Associated to Yule Processes
1. Introduction \ldots
2. A Bins and Balls Problem in Random Environment
3. Convergence of Point Processes
4. Asymptotic Behavior of the Indices of the First Empty Bins 145
5. Rare Events
6. Generalizations
7. Comparison with the Renewal Case
Appendix A. Proof of Proposition 5.1
References of Chapter IV
Comprehensive list of references

Résumé

Le but de cette thèse est de traiter quatre problèmes motivés par les réseaux de communication modernes; les outils appropriés pour résoudre ces problèmes appartiennent à la théorie des probabilités. La résolution de ces problèmes améliore la compréhension des systèmes physiques initiaux, et contribue en même temps à la théorie puisque de nouveaux résultats théoriques, intéressants en soi, sont prouvés.

Deux types de réseaux de communication sont considérés. Les réseaux mobiles sont ces réseaux où les clients se déplacent dans le réseau indépendamment du service qu'ils reçoivent; contrairement aux réseaux de files d'attente classiques, les transitions des clients ne sont pas liées aux fins de service. Dans les réseaux pair-à-pair, la distinction entre client et serveur est abolie, puisque dans ces réseaux un serveur est un ancien client qui offre le fichier après l'avoir téléchargé. Ces derniers réseaux sont particulièrement efficaces pour disséminer des fichiers gros ou populaires.

Dans les Chapitres I et II, le *comportement stationnaire* de tels réseaux est considéré. Dans chaque cas, le réseau est décrit par un processus de Markov à espace d'état discret et à temps continu, et l'on s'intéresse à son ergodicité ou au contraire à sa transience. Une spécificité de ces deux modèles est que les taux de transition des processus de Markov correspondants sont *non bornés* : dans le cas du réseau mobile du Chapitre I ceci est dû au fait que les clients bougent indépendamment les uns des autres, alors que pour le réseau pair-à-pair du Chapitre II, cela tient au fait que la capacité du système est proportionnelle au nombre de clients.

Habituellement, l'analyse de la stabilité d'un réseau stochastique se fait par l'étude des limites d'une suite de processus de Markov correctement renormalisés, appelées *limites fluides*. Cette procédure est bien adaptée pour les processus "localement additifs", i.e., les processus qui se comportent localement comme des marches aléatoires; cette propriété disparaît quand les taux de transition sont non bornés. Ces techniques sont néanmoins adaptées pour étudier la stabilité du réseau mobile du Chapitre I : utiliser des limites fluides pour étudier la stabilité de processus de Markov avec des taux de transition non bornés représente l'une des contributions de ce travail.

Le réseau pair-à-pair du Chapitre II ne se prête quant à lui pas à ces techniques, et la stabilité découle alors de l'existence d'une fonction de Lyapounov. Un autre ingrédient clef est lié à une classe spéciale de processus de branchement. Ces nouveaux processus de branchement sont définis et étudiés dans le Chapitre II, et des estimations sur leur temps d'extinction permettent, avec des arguments de couplage, d'établir des résultats de stabilité du réseau stochastique. Outre le comportement stationnaire des réseaux pair-à-pair, leur comportement transient peut aussi être étudié : ce comportement est l'objet du modèle simple du Chapitre III. Ce modèle se concentre sur l'initialisation d'un réseau pair-à-pair dans un scénario d'arrivée en masse : au temps t = 0, un pair propose un nouveau fichier que N autres pairs veulent télécharger. Contrairement au modèle du Chapitre II, ici le flot d'arrivée de nouvelles requêtes n'est pas stationnaire, il est initialement très intense puis le devient de moins en moins. Bien que le système démarre avec un serveur et beaucoup de clients, le nombre de serveurs disponibles augmente avec le temps et l'on s'intéresse au temps nécessaire pour que le réseau se mette à niveau avec la grande demande initiale. Ce problème engendre un problème de boules et d'arnes intéressant en soi, qui est traité dans le Chapitre IV.

Dans ce problème de boules et d'urnes, la distribution de probabilité qui décrit la manière dont les boules sont jetées est aléatoire : il s'agit donc d'un problème de boules et d'urnes en *environnement aléatoire*. De plus, les boules sont jetées dans un nombre infini d'urnes. Les problèmes de boules et d'urnes avec une infinité d'urnes sont bien étudiés, mais les résultats sur les problèmes de boules et d'urnes en environnement aléatoire sont peu nombreux. Quand il y a une infinité d'urnes, on peut s'intéresser à des quantités géométriques telle que l'emplacement de la première urne vide. De telles quantités ont parfois été étudiées dans des travaux antérieurs, en environnement déterministe : ici, grâce à l'utilisation de *processus ponctuels*, nous décrivons d'un coup tout le paysage des premières urnes vides, ce qui diffère des travaux précédents.

En résumé, cette thèse contribue à la modélisation des réseaux mobiles et pairà-pair; d'un point de vue technique, des problèmes liés à la stabilité des processus de Markov, aux processus de branchement et aux problèmes de boules et d'urnes sont résolus.

Summary

The goal of this thesis is to solve four problems motivated by modern communication networks; the appropriate tools to solve these problems belong to the theory of probability. Solving these problems gives insight into the original physical systems, and contributes at the same time to the theory since new theoretical results of independent interest are proved.

Two kinds of communication networks are considered. *Mobile networks* are these networks where customers perform trajectories within the network independently of the service they receive; in contrast with classical queueing networks, transitions of customers are not triggered by service completions. In *peer-to-peer networks* the distinction between clients and servers is abolished, since in these networks a server is a former client that offers a file once it has downloaded it. These last networks are especially efficient in spreading large or popular files.

In Chapters I and II, the *stationary behavior* of such networks is considered. In each case, one describes the network through a discrete state-space, continuous time Markov process, and establishes its ergodicity or transience. A specificity of these two models is that the transition rates of the corresponding Markov processes are *unbounded*: in the case of the mobile network of Chapter I this is due to the fact that customers move independently of one another, while for the peer-to-peer network of Chapter II this is because the capacity of the system is proportional to the number of customers.

Classically, to analyze the stability of a stochastic network, one can study the limits of a sequence of suitably rescaled Markov processes, the so-called *fluid limits*. This scaling is well suited for "locally additive" processes, i.e., processes which locally behave as random walks; this is however not the case when the transition rates are unbounded. These techniques are nonetheless adapted to study the stability of the mobile network of Chapter I: using fluid limits to study the stability of Markov processes with unbounded transition rates represents one of the contributions of this work.

The peer-to-peer network of Chapter II is not amenable to the same techniques, and Lyapounov type arguments are used. Another additional key ingredient is related to a special class of branching processes. These new branching processes are defined and studied in Chapter II, and estimates on their extinction time make it possible, thanks to coupling arguments, to derive stability results on the stochastic network.

In addition to the stationary behavior of peer-to-peer networks, their *transient* behavior can also be studied: this is the object of the simple model of Chapter III.

It focuses on the initialization of a peer-to-peer network under a flash crowd scenario: at time t = 0 a peer proposes a new file that N other peers are interested in downloading. In contrast to the model of Chapter II, here the flow of incoming requests is not stationary, it is initially very intense and then becomes sparser and sparser. Although the system starts with one server and many clients, as time goes by there are more and more servers available and one is interested in the time needed for the network to cope with the initial high demand. This problem triggers a *bins and balls problem* of independent interest, which is treated in Chapter IV.

In this bins and balls problem, the probability distribution that describes how balls are thrown is random: it is therefore a bins and balls problem in *random environment*. Moreover, balls are thrown in an infinite number of bins. Bins and balls problem with an infinite number of bins are well-studied, but results on bins and balls problems in random environment are scarce. When there are infinitely many bins, one can be interested in geometric quantities, such as the index of the first empty bin. Such quantities were sometimes studied in earlier works and in a deterministic environment; here, using *point processes*, we could describe at once the whole landscape of the first empty bins, which differs from previous works.

In summary this thesis contributes to the modeling of mobile and peer-to-peer networks; from a technical standpoint it solves problems related to the stability of Markov processes, to branching processes and to bins and balls problems. Since applications inevitably involve simplifying assumptions that focus on some features of a problem at the expense of others, it is advantageous to begin by thinking about simple experiments, such as tossing a coin or rolling dice, and later to see how these apparently frivolous investigations relate to important scientific questions.

— Encyclopædia Britannica, article Probability Theory

Basic research is like shooting an arrow into the air and, where it lands, painting a target.

> — Homer Burton Adkins, American chemist (1892–1949)

Introduction

Foreword

"Really, there are new things to discover in mathematics?": since I started my Ph.D., I have been faced many times in casual conversations with this simple question. Although the general public is probably convinced that mathematics is useful, this recurrent question makes me believe that they need to be convinced that current mathematical research is relevant as well. This is maybe due to the way everyone is exposed to mathematics in early education, which can give the impression that it is carved in stone.

Actually, mathematics is a very lively field, and one of the reasons for that is that it is derived from real-world problems. If a (necessarily fuzzy) line had to be drawn between pure and applied mathematics, one could probably say that pure mathematics studies abstract objects in their own right, whereas in applied mathematics one usually has a concrete motivation in mind, however remote.

Within this definition, the area of applied mathematics would for instance include the fields of optimization and of numerical analysis, as well as another field most important here, namely probability theory.

Interest in probabilistic questions arose in connection with games of chance. In the 17th century, Pascal tried to answer to the following problem asked by Chevalier de Méré: suppose two players play a certain game, whose winner is the one who wins at least four out of seven series, and are interrupted before they can finish. How should the stake be divided among them if, say, at the time they are interrupted one has won three series and the other one? It is clear that the player who has already won three series is more likely to end up winning, and should therefore get a larger share of the stake; however the possibility that the other player eventually wins cannot be completely ruled out. Intuitively, the stake should be divided according to the probability of each player of winning the game, given the current situation when the game was interrupted. Probability theory makes it possible to evaluate this probability, but only within a well-defined mathematical model.

Indeed, to answer to this question a probabilist needs to make mathematical assumptions. For instance, one can assume that the game is memoryless and fair, i.e., the results of two different series are independent and each time each player wins or loses with equal probability; one can then compute that the player who has already won three games will win with probability 7/8. However this answer heavily depends on the assumptions made, and the model can be refined to account for the fact that the players can learn from the previous series, or that one player is better than the other, etc... Although games of chance, and related fields such as insurance or finance, provide a natural framework for probability theory,

probabilistic models appear in many other settings as well, such as engineering, chemistry, biology, physics,

In the engineering for instance, one is interested in the failure of machines in supply chains, and a good understanding of these events makes it possible to efficiently dimension the supply chain, e.g., the number of spare machines. In chemistry, molecules in a medium move randomly and interact when they meet, so that the frequency of chemical reactions is governed by the concentration of each component: a probabilistic model can for instance give insight into the time needed for some rare reaction to occur, or into the time needed to exhaust all the compounds and the state of the system at that time. Similar events occur in the cells of living beings, where RNA strands and ribosomes interact when they meet RNA is then translated into proteins. In biology again, different probabilistic models can be used to shed quantitative light into the qualitative behavior predicted by the theory of evolution, considering mutations in DNA sequences as random events. Another example comes from physics: there an important problem consists in understanding how water percolates through a porous medium, such as a stone, and this simple problem has given birth to the domain of probability called percolation theory. In this setting one typically assumes that holes are randomly located within the porous medium, and the amount of water that can go through then depends on their frequency, shape, etc... Finally, a striking example comes from the study of the human brain, called brain mapping. Modern machines make it possible to know which parts of the brain respond to a given stimulation, and so a biologist is given the values of some measurements defined over the surface of the brain. However one needs to discriminate in these measurements between the signal statistically meaningful and the noise: although very abstract in the first place, probabilistic models — called Gaussian random fields — turned out to be efficient tools to answer this question.

This eclectic list of problems shows that probability theory is encountered in many different situations, which reflects the fact that randomness is an intrinsic component of many physical systems. A common feature of the above problems is that they represent a great source of inspiration for probabilists. Most of the time new questions triggered by concrete real-world problems cannot be directly dealt with existing results, and one needs to extend the theory to cover these new cases. The implications of these advances are then difficult to predict: although initially motivated by physical considerations, percolation theory is now an essential tool for the analysis of some communication networks. Moreover, the interactions between concrete problems and probability theory are two-sided: in the example of the brain mapping, the biological motivations triggered tremendous advances in the area of Gaussian random fields, which allowed in return a better understanding of concrete questions.

The job of an applied probabilist is therefore twofold: on the one hand he or she has to find relevant and interesting models to study; on the other hand he or she needs to prove rigorous mathematical results on these models. This second point is really a bottleneck, since usually, only "simple" mathematical models are tractable, e.g., models which are not too general, or with few parameters. But from a modeling point of view, the model studied needs to be representative of the physical system: to borrow a saying from one of my former professors, the goal is to study "the simplest model which is not simplistic". In other words, one is interested in studying the simplest model which both exhibits mathematical difficulties, and is at the same time relevant for the initial physical problem.

In this thesis models motivated by communication networks are considered; the main contribution is however mathematical, insofar as the results only provide a limited insight into concrete networks. Each model sets the focus on an important characteristic of the network, which makes it possible to carry on a thorough mathematical analysis.

Outline of the Introduction. The next section is intended to introduce some key concepts, related to the modeling of communication networks, to *non-experts*: due to their importance throughout this thesis, the two notions of *mathematical queues* and of *stability* are discussed from a general perspective. The goal is to give a compelling — though very subjective — rationale to explain why such questions are both interesting and challenging. It must be emphasized at this point that stability issues play an important role in this thesis, since two out of the four chapters are devoted to them.

Because of the audience targeted by this first section, the reader already familiar with queueing theory and stability of stochastic networks may wish to directly proceed to the second part of this introduction, on page 9. There the two types of communication networks that have kept me busy for three years are presented, namely *wireless* and *peer-to-peer* networks. This section gives an overview of the rich variety of problems dealt with in the existing literature, which gives a natural opportunity to introduce the models constitutive of this thesis.

The last section is devoted to the mathematical content of the different chapters. A large part of this section is concerned with stability questions, and discusses *renormalization techniques* which allow the analysis of complex stochastic networks. Other techniques are discussed as well, namely *branching processes* and *bins and balls* problems. These two models are classical tools in applied probability and are used in a wide variety of settings; here they (surprisingly) appear when studying peer-to-peer networks.

Finally, a detailed presentation of the four different chapters of this thesis follows. A reader eager to know precisely which probabilistic models are studied here is referred to this part of the introduction, which starts on page 18.

Modeling of Communication Networks

Communication networks have a huge impact on every aspect of modern societies, and constitute an interesting and stimulating domain of application for mathematics. In the Paris area, a recent survey has exhibited an average of 1.2 cell phones per inhabitant: since their appearance, mobile phones have been adopted by the vast majority of the population, representing both an important change in people's way of life and the development of a new and major economic sector. A similar observation goes with the Internet: email, e-commerce, social web sites, to name only a few, are now part of the daily life of hundreds of millions of users. In the corporate world as well, the rapid growth of the Internet has changed firms' business models, internal organization, etc...

Due to this prominent impact, communication networks are the focus of a strong research and development activity, which aims at defining more efficient networks, as well as offering new services to users. Communication networks thus constitute a lively field which evolves rapidly, and every innovation represents a new and potentially challenging mathematical problem. The scope of these problems is very large, ranging from physical considerations — how is information transmitted? — to economic ones, such as the pricing of service. Among these possible questions, the focus of this thesis is set on *performance analysis*. This encompasses the analysis of various characteristics of a network, such as the transmission time of a message, the number of customers in steady state, etc... Of course not all of these aspects are treated in this thesis, and it will become clear later on in this introduction which are.

Peer-to-peer technologies and web surfing on PDAs are among the recent technological advances that we have investigated. Although many other innovations offer interesting mathematical problems, it must be noted that every technological advance does not necessarily yield interesting mathematical questions. For instance, two different versions of the Wi-Fi protocol 802.11, such as 802.11a or 802.11b, do not necessitate two different mathematical treatments, at least for us. However, wireless Internet networks and networks of mobile phones are fundamentally different: for the former, any number of users can connect to a hotspot, at the expense of experiencing a slow connection, whereas for the latter, only a given number of users per base station can have simultaneous conversations, whose transmission rate is fixed. Such a qualitative difference between these two behaviors indeed necessitates two different mathematical models. Similarly, peer-to-peer networks offer distinctive characteristics.

Erlang. The first communication network that was the object of careful mathematical analysis is arguably the telephone. The birth of queueing theory is indeed quite unanimously attributed to Erlang, who published in 1909 a celebrated paper [**Erl09**], "The theory of probabilities and telephone conversations". Together with his second seminal paper [**Erl17**] published in 1917, Erlang investigates the problem of *dimensioning* a telephone network.

Basically, a telephone network is a collection of telephone lines, to which users access and which they keep busy for the duration of their conversation. The same description applies to more modern networks of cellular phones, where it is a common experience to initiate a call that is rejected because the network is busy; in other words, all lines are occupied, and one must try again later and hope that a conversation has ended in the meantime. Setting asides physical considerations such as the quality of the communication itself, the *quality of service* perceived by the customers is measured through the frequency at which this event happens, which is called the *blocking probability*; a good quality of service means a low blocking probability. The primary wiggle room for a network operator is the total number of available lines: more lines mean a better quality of service but obviously a more expensive network as well. Thus the operator needs to tune the trade-off between the quality of service perceived by the users and the cost of the network.

Although evaluating the cost of a network is probably not a difficult issue (by which we mean that it does not require sophisticated mathematical reasoning), the blocking probability represents on the other hand a more challenging task. It is important to note that given the network, the blocking probability is solely determined by the customers' behavior, i.e., the frequency and the duration of their calls. As an illustrative example, imagine that one wants to dimension a network for 1,000 customers who make, on average, 4 calls a day of mean duration 5 minutes. The number of customers being large, the law of large numbers predicts that random fluctuations around the mean values will average out, so that finally, around 20,000 minutes of communication will be generated during a typical day. One single line providing 3,600 minutes of communication a day, a rule of thumb advocates the use of 6 telephone lines. However, this result does not say anything about the blocking probability, and for a good reason: the law of large numbers tells about the long term behavior of the network, while a call is rejected when a customer initiates a call and all lines are busy. Thus the blocking probability is determined by the instantaneous state of the network, which cannot be taken into account by law of large numbers type of arguments; it is intrinsically of a different nature. These events largely depend on the random behavior of the customers: thus as indicated in the title of Erlang's first paper, probability theory is the natural framework in which to cast this problem.

The problem can then be formulated as follows: given customers who originate calls at random times, each call being itself of a random duration, what is the probability that a call finds all lines busy? A fundamental and natural assumption is to assume that customers behave independently of one another: under this assumption, Erlang argues that calls arrive in the network at the epochs of a Poisson process. As for the call duration, fields measurements performed by Erlang show that they follow the exponential distribution. These different assumptions define completely the problem from a mathematical standpoint, which Erlang solves to prove the famous Erlang B formula. It gives a flavor of the kind of results encountered in queueing theory: if $\rho > 0$ is the traffic intensity — equal to $\lambda \times r$ where λ is the mean number of calls arriving in a unit of time and r is the mean duration of a call, measured in the same units — then the blocking probability $P(\rho, n)$ when there is a total of n lines is given by

$$P(\rho, n) = \frac{\rho^{n}/n!}{1 + \rho/1! + \dots + \rho^{n}/n!}$$

Plugging in the values of the above simple example, one finds a blocking probability of 23% with six lines, so that the random fluctuations can hardly be neglected. Conversely, this formula can be inverted to find that twelve lines are needed to ensure a blocking probability of less than 1%: having a closed form formula allows such manipulations, which is one of the strength of this theoretical approach. The Erlang B formula was moreover shown to hold even if call durations are not distributed according to the exponential distribution, see for instance Takács [**Tak69**]. Although traffic patterns in telephone networks have changed drastically since Erlang, the robustness of this formula probably explains its enduring success; it is still used by today's network operators.

Queues. In a telephone network, customers finding all lines busy will retry later on. Another possibility would be to queue these customers, so that the earliest customer that arrived and found all lines busy could use the first line to get free; this mechanism describes for instance the operation of a call center, or of a supermarket with several cashiers and a single queue. These two simple examples illustrate the concept of *mathematical queue*, and already show that a single mathematical model can be representative of different physical systems.

In a supermarket, customers are served based on their order of arrival, and we say that the service discipline is FIFO or FCFS, for First-In First-Out and First-Come First-Serve, respectively. Other service disciplines could make sense: if one knew the service required by a customer, i.e., the amount of time this customer will keep a cashier busy, then it could be efficient to determine the order in which to serve customers based on this information — e.g., by treating preferentially customers with only few items. Last-In First-Out (LIFO) is another example of a simple service discipline, it models for instance a stack of items, which is a common object, for instance in computer science. When a new item arrives, it is placed on top of the stack, so that the first item to be processed is the last one to have arrived. Similarly, the processor of a computer can be seen as a queue with a special service discipline: the processor divides time into slots, devotes its computing power during one slot to one application, and successively inspects the running applications slot after slot. This gives the user the impression of parallelism and indeed enables him to launch several applications simultaneously. Each application creates an incoming flow of jobs that the processor needs to handle, and each job necessitates more or less computation from the processor. When the length of a time slot is very small, a reasonable modeling assumption amounts to suppose that jobs are treated simultaneously, and so the service discipline is not FIFO, it is a well-known discipline called Processor-Sharing.

These different examples illustrate the different parts that define a queue from a mathematical standpoint. First there is indeed a queue — e.g., the customers placed on hold in a call center or the jobs waiting to be treated by a processor. Then there are *servers*, which serve the customers — e.g., the call operators or the processor. And finally there is a service discipline, which determines in which order customers are served — e.g., FIFO or Processor-Sharing. Finally, to fully define a mathematical queue, one needs to specify the statistical hypotheses concerning the customers' behavior, namely call arrivals and durations. In the queueing terminology, the duration of a call is often called the *service requirement*, because this corresponds to the amount of time it keeps a server busy.

Queues are fundamental objects to model communication networks. The example of call centers has already been given, but the Internet provides a much more important and richer example. The Internet can indeed be seen as a network of FIFO queues, which gives a good idea of modeling of communication networks. Due to the Transfer Control Protocol (TCP), which plays a fundamental role in the stability of the Internet, it can also be partially modeled as a network of Processor-Sharing queues, but this is a more delicate story.

The Internet is the network of computers that communicate using the Internet Protocol (IP), which, broadly speaking, defines the address of a computer, the famous IP address. It is not to be mistaken with the Web, which is the network of web pages, and as such a "subnetwork" of the Internet; for instance, peer-to-peer networks use the Internet but have nothing to do with the Web. To communicate over the Internet, two computers send *packets* to each other: these packets are relayed from the source to the destination through a series of dedicated machines, called *routers*. The route — i.e., the sequence of routers — taken by a packet is not known in advance: each router has a buffer, where incoming packets are stored. The router inspects the packets buffered sequentially, and decides, based on the

destination IP address, the router to which to forward the packet. If a router is seen as a server, its buffer as a queue and packets as customers, then the Internet can naturally be seen as a network of FIFO queues, where a customer goes from one queue to another until it exits the network; under particular statistical hypotheses, such networks are called *Jackson networks*.

The above examples show that mathematical queues provide a flexible framework to model a great variety of physical systems. Yet because of this versatility, queueing theory lacks of a unified theoretical framework. When one gets interested in a new class of network, it is illusory to think that it will readily fall within the range of some existing theorem, as well as it seems illusory, for now, to try to formulate and prove a theorem that would encompass many different networks. So the corpus of queueing theory consists essentially (although not only) in results about various network models, and this partly explains the global structure of this document, where different chapters essentially correspond to independent problems. It seems all the more hard to unify different results under general theorems that even two close models can exhibit radically different behaviors.

For instance, although the behavior of two queues in tandem is well understood — this can model the waiting time at some furniture shops, where one first has to queue at the cashier, and then a second time to retrieve the item — Bramson [**Bra94a**] studies a close model where customers can re-enter the network after being served the second time. This slight difference yields a different and counterintuitive behavior, related to the fundamental notion of *stability*.

Stability. The simplest way to think of stability is the following: imagine a network is a black box that transforms an input flow into an output flow. Clearly the output flow cannot be larger than the input one — at least in the long run — and we say that the network is stable if the output equals the input. Otherwise, the output is strictly smaller than the input and some flow necessarily accumulates in the network, which is then said to be unstable. Studying the stability properties of a network can therefore be thought of as studying its *capacity*, i.e., the maximal input rate it can accommodate; we discuss below why the situation is (un)fortunately sometimes not that simple, and in order to illustrate the subtlety of this notion, we spend some time discussing some surprising results.

But before that, note that the notion of stability is not always appropriate: since the total number of customers in a telephone network is upper bounded by the number of lines, customers cannot accumulate indefinitely and so it is always stable. As discussed before, the relevant question in that case concerns the blocking probability. In contrast, in the previous simple model of the Internet, packets can accumulate in the buffers of the different routers, and the network can then be stable or unstable. To each case corresponds different questions: Under stability conditions, one wants to know more about the steady state, for instance the average number of packets that remain in the network or its fluctuations in time; when it is unstable, it can be interesting to characterize the rate at which packets accumulate, or where and how in the network do these packets aggregate.

Stability properties of the single-server queue have been understood in a fairly general setting since 1962 with Loynes [Loy62]. If the server works at rate μ , then the queue is stable if and only if the arrival rate λ is smaller than μ , i.e., $\lambda < \mu$.

In other words, the queue is stable if and only if the server can output more fluid than what enters. Under stability conditions, although the instantaneous output rate is $\mu > \lambda$ when the server is busy, the long-term output rate from the queue is exactly equal to λ due to idling periods of the server. The condition $\lambda < \mu$ is often rewritten $\rho < 1$, where one defines $\rho = \lambda/\mu$ as the traffic intensity.

For networks of queues, one can still define the traffic intensity: the input rate to a queue is for instance given by the rate of exogenous arrivals plus the rate of arrivals due to inner transitions of customers, which themselves depend on the arrival rates into the other queues; thus the arrival rates are usually determined through fixed point equations. In view of Loynes' results, a natural idea is to think that a network of queues will be stable if and only if the traffic intensity at each node is smaller than one. One way is correct: if the traffic intensity at some node is greater than one, then the network is unstable, since customers will necessarily accumulate at this node. The converse has been the object of intensive investigations in the early '90s. Because of Loynes' results, the simplest setting where this conjecture could fail is when the network consists of at least two queues. And indeed, Rybko and Stolyar [RS92] in 1992, and Bramson [Bra94a] two years later, came up with counter-examples for this yet appealing conjecture with networks consisting of only two queues. In both examples, the customers accumulate within the network as follows: one of the two queues is large and the other empty, and then the large queue empties while the empty queue builds up. Although each queue is empty infinitely often, the amplitude of the oscillations are larger and larger.

Such situations naturally appear in *multiclass* queueing networks, where customers are of one out of several possible types, or classes. Different classes (may) differ by their arrival processes or the routes of their customers within the network, and more importantly because servers can prioritize customers based on their classes. Rybko and Stolyar's counter-example [**RS92**] indeed relies on a two-node queueing network with two classes of customers, where each server gives strict priority to one class of customer over the other. This results in an unstable network which nonetheless satisfies the condition $\rho < 1$ at each server. In a similar vein, Dumas [**Dum97**] exhibits a three-node network with two classes of customers where the stability region is non-linear, non-convex and non-monotonous. This means in particular that the network may be unstable for some arrival rate λ and still stable for some larger arrival rate $\lambda' > \lambda$. Thus the concept of maximal input rate that a network can accommodate does not always make sense.

In each of the two above counter-examples, servers give strict priority to one class of customers, which can seem artificial. To compensate for this unsatisfactory situation, Bramson [**Bra94a**] built an unstable two-node FIFO queueing network with only one class of customers where the usual conditions $\rho < 1$ are satisfied. And still in 1994, Bramson exhibited an even more striking behavior in [**Bra94b**]: He built a class of networks such that if the traffic intensity ρ at each node is smaller than some threshold $\rho^* < 1$, and then it may be *arbitrarily small*, then the corresponding networks are unstable.

In conclusion, the question of stability represents both an important and challenging issue. Although the above various counter-examples highlight some possible subtleties, there are nonetheless many networks where the usual conditions $\rho < 1$ give the correct stability region; this is for instance the case for Jackson networks, mentioned previously. This is also the case for the network models investigated in this thesis, which we now introduce.

Wireless and Peer-to-Peer Networks

Two types of networks are studied in this document: wireless and peer-to-peer networks. The goal of this section is to give a broad overview of the issues raised by these networks, and to explain the models and questions studied in related works. In doing so, the models studied therein are naturally introduced, and the issues addressed and models considered are positioned. As mentioned in the beginning of this introduction, the reader is referred to the next section for details on the technical (i.e., mathematical) content of this thesis.

We first introduce wireless networks, and then peer-to-peer networks. Since the latter are built on top of the Internet, a quick overview of the wide range of problems raised by the Internet is done before introducing peer-to-peer networks. As for wireless networks, the two different types of wireless networks, namely *ad-hoc* and *infrastructure* wireless networks, are introduced. It must be stressed at this point that this thesis does not contribute to the field of ad-hoc wireless networks; nevertheless, the range of issues raised by such networks is wider, in my opinion, than the issues raised by infrastructure wireless networks. For this reason, time is spent in introducing some interesting issues specific to ad-hoc wireless networks. Note that although ad-hoc wireless networks seem richer from a mathematical standpoint, the vast majority of real-life networks have a fixed infrastructure.

Wireless Networks. Plainly, these are networks where communication between two nodes is carried by electromagnetic waves, propagated in the air. We sometimes use the generic term of node, since as will be seen, users can communicate directly with one another, but communication can also occur between a user and a base station; a node can be a sensor too, in the case of sensor networks. Because they exhibit original features compared to wired networks, such as the Internet or the telephone, wireless networks have triggered a large amount of work in various fields, such as information theory, algorithmic, random geometry or scheduling.

Ad-Hoc Wireless Networks. These are wireless networks where any two nodes can directly communicate with one another, without resorting to a fixed infrastructure. Since information is transmitted through the air where the signal fades away, communication between two nodes is possible only if they are within transmission range, i.e., close enough. In multihop networks — a hop informally refers to the transmission of a message from one node to another — nodes that are further apart can communicate by using other nodes as relay. The message is then forwarded from node to node until it reaches its destination. It is usually not possible to restrict the direction in which the radio signal is emitted, so that the drawback of allowing any two nodes to directly communicate is that, when they do so, they create a zone of interference where other nodes' communication is hindered. Two models for this interference zone are equally prevalent in the literature, and most of the following references consider both of them. On the one hand, one can consider that every communication is strictly prohibited in this zone. Then as soon as a node receives a signal from another node, it does not emit itself; because the situation is binary — a node either receives a signal or not — such models are called *boolean models*. It is usually assumed that any node has a fixed transmission range, thus prohibiting communications in a given radius: the interactions are short-range. On the other hand, in the *physical model*, a node can separate the signal it is interested in from the other signals it is exposed to — the noise — as long as the *power* of the signal of interest is larger than the power of the noise: the *signal-to-noise ratio* needs to be larger than some threshold for the communication to be successful. It is usually assumed that the power decays smoothly with the distance, so that any node influences every other node, however weakly: the interactions are long-range. In this setting, power control is an important issue: it determines the range of a signal as well as the strength of the interference created. Power control is moreover a critical issue for autonomous devices, since it determines their life span.

In either model, interference has an adverse impact on the network's capacity, which is defined as the maximal rate at which nodes can transmit messages. Imagine for instance that each node has a buffer, where messages waiting to be transmitted are stored. If new messages appear at a rate higher than the network's capacity, then the number of queued messages will become larger and larger, and the network will be unstable. With a low density of nodes, interference does not play a significant role, and all communications will essentially be accepted. But as the density increases, more and more communications will be inhibited and an interesting question is to quantify this impact. Information theory sheds light on these questions by providing theoretical upper bounds on the network's capacity; an algorithmic issue then consists in designing algorithms that reach, or get close to, this theoretical upper bound. For instance, Gupta and Kumar [GK00] have shown in 2000 that if n nodes are arbitrarily located in a network of fixed area, then the throughput of any node vanishes like $1/\sqrt{n}$ as n gets large; they moreover exhibited a deterministic scenario where this bound is reached. This scenario is nonetheless unrealistic — nodes need to be regularly spaced — and they complement their study by looking at the situation where nodes are randomly spread over the network. They found that the capacity of each node vanishes like $1/\sqrt{n\log n}$ in this case. Although information theory does not preclude the capacity of reaching $1/\sqrt{n}$, the additional factor $1/\sqrt{\log n}$ was seen as the price to pay for randomness.

The main point in the case of randomly located nodes is to define a routing algorithm, i.e., an algorithm that determines how to forward a message from the source to the destination. The routing algorithm originally proposed by Gupta and Kumar [**GK00**] tries to connect two nodes by a straight line, and its analysis relies on random geometry. Seven years later, Franceschetti *et al.* [**FDTT07**] showed that the factor $1/\sqrt{\log n}$ was actually due to this particular choice of routing algorithm, and proposed another routing algorithm that reaches the information theoretic bound $1/\sqrt{n}$ in the case of randomly located nodes. This new algorithm is based on the existence of a particular configuration of nodes. Namely, Franceschetti *et al.* show that the network can organize a highway of nodes, in charge of relaying the long-distance information transmitted over the network.

This example illustrates the interplay between the *connectivity* properties of the network and its capacity. The right mathematical tool for formalizing this interplay is the *percolation theory*, which is indeed the cornerstone of the arguments of Franceschetti *et al.* [FDTT07]. See for instance the first chapter of Grimmett [Gri99] for an introduction to this theory, where some physical motivation is given. This mathematical framework turned out to be very fruitful in the context of wireless networks. In addition to the aforementioned results of Franceschetti *et al.*, Dousse *et al.* [DFT06] give an explanation to the fact that the throughput seen by any node vanishes as the number of nodes increases: using percolation theory, they prove that this is actually the price to pay to get the full connectivity of the network. If a fraction of nodes cannot communicate with each other, then the network can be designed in such a way that the throughput does not vanish in the limit.

The above mentioned results concern ad-hoc networks where users do not move, and a natural extension consists in allowing mobility. In this case the connectivity properties of the network evolve over time, so that the results in the fixed setting need to be revisited. A first difficulty consists in designing efficient routing algorithms, see, e.g., Tschopp *et al.* [**TDG08**] and the references therein. An interesting discovery of Grossglauser and Tse [**GT01**] is that the situation under mobility can be much more favorable than in the fixed setting: although in the fixed setting the network's capacity decreases like $1/\sqrt{n}$, they show that the throughput can be kept constant if users move. The underlying idea is that mobility represents an opportunity for the network to increase its capacity by devoting its resources to nodes that are in a good state; we will shortly come back to this interesting property in the context of infrastructure wireless networks.

Infrastructure Wireless Networks. Ad-hoc wireless networks are those wireless networks which can operate without fixed infrastructure. In wireless networks with a fixed infrastructure, be it a hotspot in the case of the Wi-Fi or a base station in the case of cell phones, communication is always between a user and this infrastructure, say a base station for simplicity; such networks are sometimes called *infrastructure wireless networks*, a terminology that we use here. This communication is directed, from the users to the base station (many-to-one), or from the base station to the users (one-to-many), but we will not enter this level of details. Since several users can connect simultaneously to a base station, this latter needs to divide its capacity according to some *scheduling policy*; these problems are sometimes referred to as *bandwidth-sharing* problems.

In the case of cell phones, the data transmitted is the voice, which is sent at a fixed rate determined by the data compression scheme used. In particular, users require the same transmission rate — the traffic is said to be *inelastic* — so there is a simple way for the base station to divide its bandwidth. Imagine for instance that every user requires 13 kbit/s, and that the base station's capacity is equal to 100 kbit/s: then the base station will divide its capacity into 7 slots of equal capacity of 13 kbit/s, thus being able to serve a maximum of 7 users simultaneously. In the case of the Wi-Fi, this question is more complex, since now users are happier if they can get a faster connection. A simple approach is to serve all the users simultaneously and equally, i.e., if there are 10 users connected, then each user would receive one tenth of the base station's capacity. This answer is satisfactory in a fixed setting, when users' capacity — the amount of data they can send or receive in a second — does not evolve over time. Knopp and Humblet [KH95] investigate the situation in a dynamic setting, where users have a time-varying capacity. They show that under power constraints, the base station should at any time serve only one user, the one with the best capacity, in order to maximize the network's throughput. This scheduling induces delay in message transmission, because a user will not be served until its state is favorable, and so the information relayed needs to be delay-tolerant. Users have similar behavior, so that any user will indeed be served after some time, but in general, one must be careful that such policies are fair and do not induce starvation of a category of users.

The case of time-varying capacity is not just an intellectual game, it naturally arises in the context of mobile users. The capacity of a user connected to a base station indeed depends on how the user "sees" the base station, determined by the distance between the base station and the user, the presence of obstacles between them, etc... Thus under mobility assumption, users' capacity will naturally vary over time. Knopp and Humblet's results suggest that the optimal scheduling policy then consists in serving the user with the best capacity. This observation motivated the above mentioned result of Grossglauser and Tse [GT01] in the case of ad-hoc wireless networks, while the situation has been extensively investigated in the case of a single-cell network with a base station, see for instance Bonald *et al.* [BBP04] and the references therein. Knowing the users' capacity implicitly assumes that the base station is aware of the channel conditions, and these scheduling disciplines are called *channel aware* scheduling disciplines; the above mentioned works of Bonald et al. [BBP04] and other results have shown that such algorithms improve significantly the throughput. In practice a network consists of more than one base station: with several stations, new problems appear due to the fact that customers move from one cell to another, and also because neighboring cells can interfere. Interference problems are not discussed here.

The situation then depends on whether the traffic is elastic or inelastic: for inelastic traffic such as the voice, the capacity of each cell is finite, i.e., there is a maximum number of users per cell. A user moving to a cell with already the maximum of users will see its service interrupted, and such events need to be controlled since they are worse than a call simply being rejected initially. When several types of users coexist, which may typically correspond to different traffic patterns, Antunes *et al.* [AFRT06, AFRT08] exhibit a stability problem: the system spends a long time in a certain state — e.g., favoring a certain class of users — and then switches for another state, where again it stays for a long time. This is a bad property from an operator's perspective, since one wants to be able to guarantee some quality of service.

For elastic traffic, any number of users can be simultaneously in a cell: again problems of interference and power control are important, but also the problem of bandwidth sharing. The base station needs to decide which users to serve, and how to divide its capacity. This question, investigated in depth for the single cell, was extended by Borst *et al.* [**BPH06**] by considering inter-cell mobility: the model of this paper is of primary interest for us. This paper considers a network of base stations, where users enter the network with some service requirement, and then move within the network independently of the service they receive until their initial service requirement has been met. Along their route, they share the capacity of the cell they are in with the other customers simultaneously present. This model is very flexible: Borst *et al.* consider different classes of customers (corresponding to different service requirements and routes), each cell is divided into regions corresponding to different radio conditions, and base stations implement sophisticated scheduling disciplines. Their primary objective is to characterize the stability region of the network for different scheduling disciplines. They use *fluid limit* techniques to identify this region — these are discussed in the next section — but their analysis is not completely rigorous. A serious technical difficulty arises because of the coexistence of two different time scales, one for the variation of the total number of customers in the network, and the other for the movements of customers within the network.

Chapter I of this thesis is devoted to fixing this problem. In order to do so we had to develop a technical approach completely different from theirs, which in the end rigorously justifies the use of fluid limits to characterize the stability region. The model of this chapter is simpler than the one of Borst *et al.* insofar as there is only one class of customers, and more importantly the base stations implement the simplest possible scheduling discipline. The conclusion of this analysis is that thanks to the users mobility, the stability region is as large as it can be; similarly as in the above mentioned results, the mobility has a positive impact on the network's capacity. Note that although the model studied in this chapter finds its motivation in the modeling of infrastructure wireless network, its contribution is in the end essentially technical, as is explained in the next section.

In contrast to the model studied in Chapter I which contributes to the field of wireless networks, Chapters II and III are motivated by the modeling of peer-topeer networks. These are particular networks built on top of the Internet, which we introduce first before going into the details of peer-to-peer networks.

Internet. Although any two computers communicate over the Internet using IP addresses, the medium as well as the content of the communication are varied. The medium can be the air, optical fibers, ADSL or even electric lines, whereas the communication taking place over this versatile medium can be file sharing, telephony, parallel computing, etc... In the Web, a communication takes place between a user and a web server, which hosts the web site of interest, using the HTTP protocol. More generally, every type of communication is usually associated with one or more protocols, e.g., web browsing with HTTP and HTTPS, file sharing with FTP and BitTorrent, telephony with Skype, etc...

These many different mediums of communication each offer specificities worth investigating. To name only a few, we have already discussed in length the case of the wireless Internet as a special case of wireless networks; due to their physical nature, optical communications raise new problems, especially in routing [**ZJM00**]; ADSL lines carrying domestic traffic exhibit peculiar patterns due to peer-to-peer

or video streaming traffic, and their analysis is a subject of utter interest for network operators since they carry a majority of the current Internet traffic [**bAGP**+**05**]. In addition to these problems, the many different protocols each trigger new problems as well. BitTorrent and peer-to-peer protocols in general have been intensively studied in the past years (see below); TCP is the cornerstone that makes the Internet work, but its modeling is a challenging question that has yielded different studies [**DGR02, CMP09**]; widely used access protocols such as Aloha or Ethernet present mathematical peculiarities, such as the instability of Ethernet [**Ald87**], and are still intensively investigated currently, see the recent work of Bordenave *et al.* [**BMP08**] on the stability of Aloha; some attempts have been made to define efficient admission control protocols at the router level, such as RED [**FJ93**]. And more generally, many aspects of the Internet can be optimized or enhanced. Just to name a few, Appenzeller *et al.* [**BKM04**] try to dimension the size of the routers' buffer, and Benameur *et al.* [**BFOBR02**] are interested in introducing admission control and quality of service in the Internet.

The modeling of the whole Internet at once is a tantalizing task that reveals many challenges; some respectable researchers think that the behavior of TCP is well understood only for networks with one node. The interaction between several nodes indeed induces a high degree of complexity. Massoulié and Roberts [MR99] recently introduced a popular model for elastic flows, which sees TCP as a black box which divides the bandwidth of each router among the different flows that traverse it according to some optimization problem. This model has found some justification thanks to a recent work of Walton [Wal09], which shows that this bandwidth sharing model can be obtained by a scaling procedure starting from a rather general class of networks. Let us close this long and eclectic list of problems related to the Internet by mentioning a recent and innovative work of Bonald *et al.* [BFP09], in which they question the very need of a control protocol such as TCP, and suggest that an efficient use of source coding could do the job.

This list of problems is very far from being exhaustive — for instance, we did not discuss insensibility issues or graph-related problems such as the structure of the Internet or the PageRank algorithm — but just gives an idea of the wide range of problems triggered by the Internet. Every level of the Internet, from the physical to the application layer, raises many problems. Among these problems, we have decided to focus on peer-to-peer networks.

Peer-to-Peer Networks. A typical situation in the Web is that of a communication between a user and a web server, the former requesting a file, e.g., the HTML code source of a web page, from the latter. If several users connect simultaneously to the same server, then it needs to divide its capacity between the different users. The most common sharing policy is the Processor-Sharing discipline, where every user receives only a small fraction of the server's capacity when many of them are connected. This is not a problem if the file requested is small — a web page for instance — but if a user is interested in downloading a larger file, it may then experience a large delay. Another observation is that since the server is the only source where the file is available, the time needed to download it is essentially proportional to its size. For these two reasons, the Web and the HTTP protocol are not suited for sharing large or popular contents, such as, respectively, a Linux distribution or a movie. Although some protocols were specifically developed to address this issue, for instance FTP, the Web was still widely used at the end of the 90's to exchange large files.

Peer-to-peer networks have been designed to provide an alternative way of sharing large or popular files. The simple yet fundamental idea of peer-to-peer networks is the following: once a user has received a file, he offers it in turn, thus acting like a server. If many users want the file, many users will offer it after some time, and so it seems reasonable to expect that such networks will cope with a high demand. Note in particular that in contrast to classical web architecture, under certain conditions, the higher the demand and the better the network's capacity. The vast success of these technologies is reflected in recent figures that show that a large portion of the Internet traffic originates from peer-to-peer applications. These figures are however to be taken with some care, since measuring peer-to-peer traffic proves to be a difficult issue, see Saddi and Guillemin [SG07].

Another crucial idea that dramatically speeds up the dissemination of a large file consists in splitting it into small pieces, called *chunks*; a movie will typically be cut into a couple of thousands of chunks. This strategy has numerous advantages: first of all, the user can download different chunks simultaneously and from different sources, thus both speeding up the downloading process and balancing the load over different peers. Another key advantage is that a peer can start sharing the file as soon as he has a chunk: in particular, he can quickly participate in increasing the network's capacity, instead of having to wait to have the whole file in his possession. This idea is key in BitTorrent, one of the most popular peer-to-peer networks nowadays. The drawback is that it increases the complexity of the algorithm, and the system's performance is very sensitive to the policy that rules out which chunks are exchanged between two peers. More generally, peer-to-peer networks are challenging networks to operate, since now the information is no longer centralized but disseminated all across the Internet.

Peer-to-peer networks have to provide peers information on how to find chunks they are interested in. Not only are these chunks spread over the entire network, but the topology of the network evolves dynamically over time, when peers enter or leave the network, propose a new content or get some old one. Providing the relevant information to the peers is thus a difficult issue, and largely determines the network's (in)efficiency.

One approach consists in trying to maintain a fixed structure on the network's topology, e.g., organize the peers into a ring. One can then exploit properties of this topology to optimize the lookup of information. Some attempts in this direction consider de Bruijn graphs for the overlay topology, and rely on the use of distributed hashtables, see for instance Gai and Viennot [**GV04**] and the references therein. Due to the dynamically evolving topology of the network, maintaining a fixed structure turns out to be a difficult issue in practice. Instead, recent algorithms such as BitTorrent rely on random graph. Broadly speaking, a peer in BitTorrent connects to a random set of peers, from whom he will try to obtain interesting chunks. The set of peers a peer communicates with then slowly evolves over time, see [**Bit**] for more details.

Once peers have decided to communicate with one another and exchange chunks, another crucial question concerns the chunks that should be exchanged. The policy chosen in this respect has a great impact on the downloading time of the full file, as well as on the scarcity of chunks in the network. Because peers leave the network after some time once they have completed the download, it can indeed happen that some chunks disappear, or are rarer than others.

A related problem concerns the time needed to download the first and last chunks. It has been argued, e.g., in Tian *et al.* **[TWN06]**, that peers spend most of their time downloading their first and last chunks. For the first chunks, this comes from the fact that initially, a peer does not have chunks to exchange, and thus other peers do not have a strong incentive to communicate with him. On the other hand, the last chunk problem comes from the fact that getting one precise chunk can be difficult. To circumvent this problem, solutions implying network coding have been proposed **[GMR06]**, as well as exchange policies that favor the exchange of rare chunks, see for instance Bharambe *et al.* **[BHP06]**.

We now introduce models closely related to the ones studied in Chapters II and III. The discussion is twofold, depending on whether the system is studied in stationarity, or if the focus is rather set on the initialization of the network.

Performance Analysis in Stationarity. The key idea in peer-to-peer networks is that each peer participates in increasing the network's capacity: not surprisingly, this powerful and generic idea has spread beyond the limited scope of file sharing. Skype is a good example of the use of a peer-to-peer network in a real-time context, and such applications are an active topic of research. Recently, Bonald et al. [BMM⁺08] and Massoulié and Twigg [MT08] have investigated the performance of a peer-to-peer system used to broadcast real-time data, such as a television show. The case of live streaming presents new constraints compared to file sharing: it is not mandatory that each peer receives exactly every packet of the stream, but the delay is critical. In $[BMM^+08]$, Bonald *et al.* compare different chunk exchange policies and analyze their impact on the system's performance, while Massoulié and Twigg [MT08] establish the network's capacity under various scenarios. One limitation of these works is that they consider a static context, i.e., the network's topology is fixed and does not evolve in time. One of the motivation of our own works on peer-to-peer systems was to try to analyze in some way the impact of the evolving topology. The price to pay is twofold: the multichunk scenario is much more difficult to analyze, and it is hard to consider a realistic topology for the overlay network.

In addition to the above live streaming scenario, file-sharing peer-to-peer networks have been extensively studied: these are complex networks where several aspects can be considered. A popular model introduced by Massoulié and Vojnović [**MV05**] neglects the time needed to download a particular chunk, and puts the emphasis on the process of contacts between peers. In these models, at random times, each peer contacts another randomly selected peer and tries to exchange some chunks with him. It is assumed that once a contact is made, the peer takes a chunk uniformly at random among the chunks that the contacted peer has and that the initial peer does not have; the exchange is one-sided. Kesidis *et al.* [**KKS09**] consider a model with two-sided exchanges upon contact. In both cases, the initial stochastic system is studied by means of a deterministic approximation involving a system of differential equations, and the stability of this system is the primary concern. In contrast to the aforementioned papers on live streaming, in these models the overlay topology is complete, i.e., any peer can contact any other peer.

If chunks are indeed small in practice, it is however not clear that their downloading time can be completely neglected. Qiu and Srikant [QS04] investigate a simple queueing model for BitTorrent-like peer-to-peer networks. Their model consists of two queues in tandem, where peers in the first queue do not have the file, and peers in the second queue have it. The peer-to-peer dynamics then dictates that the first queue behaves as a queue with a varying and random number of servers, where a server is precisely a peer in the second queue. This model again implicitly assumes that there is a complete overlay topology. Because peers who do not have the file are impatient, the queueing system they propose is always stable, but when this assumption is removed, such as in Susitaival et al. [SAV06], then the stability needs to be studied. The stability criterion involves a comparison between the input rate and the mean number of peers in the second queue, which determines the mean output rate of the system. Susitaival et al. derive such a stability condition from a heuristic standpoint, and this intuitive result is proved in Chapter II under more general hypotheses. A remarkable feature of these networks is that when peers who have the file are patient enough (i.e., stay long enough once they have downloaded the file), then the network is always stable and can accommodate any input rate.

These two models do not take into account the multichunk situation, and there is a good reason for that. Although with a contact process perspective, the combinatorial problems induced by several chunks can be handled — with n chunks, each peer is characterized by one of the 2^n possible subsets of chunks — the situation is much more challenging with a queueing perspective. Indeed, the subset of chunks then not only characterizes the peers, but the servers as well. When a peer has k chunks, then its capacity, as a server, needs to be divided in some way among these k chunks. To the best of my knowledge, there is no good queueing model that correctly handles this difficulty. A possibility to circumvent it is to restrict the sets of peers that can interact. For instance, Parvez et al. [PWMC08] consider the possibility to parse a file while downloading it. The chunks of the file then need to be downloaded in order, what reduces the dimensionality of the problem since now a peer can be characterized by the number of chunks it has. A further simplification amounts to impose that a peer with k chunks always asks the next chunk from peers with k+1 chunks; in particular, a peer is a server for only one chunk, the last one it has downloaded. This architecture can be justified by load balancing arguments: without this constraint, peers with the whole file would receive more requests than other peers, and more generally the more chunks a peer would have, the more requests it would receive. These two assumptions lead to a simple network model with n+1 queues in tandem, if n is the number of chunks, and where the dynamics between two successive queues is similar to the model of Susitaival et al. [SAV06]. It can therefore be seen as a generalization of their model, that is analyzed in some particular situation in Chapter II. One would ideally like to study the case where chunks have identical size, but, among other results, we prove that if chunks are smaller and smaller (i.e., a peer starts by downloading the biggest chunk, etc...) and if customers are patient enough, then a similar phenomenon as in the single-chunk case holds, i.e., the system is stable for any input rate, which is in sharp contrast to classical queueing networks.

Flash Crowd. All the above models consider peer-to-peer systems in a stationary regime, with a continuous flow of incoming peers (or chunks in the case of live streaming), for which the problem of stability makes sense. Another opposing situation is the *flash crowd* phenomenon: this corresponds to the initialization of a peer-to-peer network upon the release of a popular file. If many peers were awaiting this release, then there will be a sudden burst of incoming peers shortly following the release. This phenomenon has been observed in real-world networks, see for instance Pouwelse *et al.* [**PGES05**] for measurements of the flash crowd effect corresponding to the "release" of the movie "Lord of the Rings III".

As long as the flash crowd effect prevails, the system is in a transient regime. An interesting question concerns the time needed to reach the stationary regime, or, put otherwise, the time needed for the system to cope with the initial high demand. Yang and de Veciana [YdV06] look at this question, but their analysis remains a first-order one, and basically amounts to say that the number of servers grows exponentially. Indeed, the system begins with one server (the one initially offering the file), which is replicated after the first customer finishes downloading the file; at this time, the system consists of two servers working in parallel, so the next server is created twice faster, etc... This dynamics is analog to the dynamics of a population of cells, where each cell splits into two identical cells after some time. This analogy with this kind of biological processes, called branching processes, makes it possible to carry out a more detailed analysis of the transient phase of the peer-to-peer systems, which is the object of Chapters III and IV. This analysis makes it possible to get a deep understanding of the system's behavior and, as a simple consequence, to justify the first-order approach used by Yang and de Veciana [YdV06].

So far, we have not entered the technical content of our works on mobile networks or peer-to-peer systems, in order to focus on the modeling problems; the next section introduces it.

Mathematical Framework

In this section, an overview of the mathematical tools used is given. All the works contained in this thesis were motivated by the modeling of communication networks, and the choice has always been made to use simple Markovian assumptions: the section starts with a few words on this topic. Since two chapters are devoted to the study of the stability of Markov processes, we then present renormalization techniques which are modern tools designed to address such issues. Finally, we conclude this section with the introduction of two common probabilistic models which naturally appeared along way, namely branching processes and bins and balls problems.

Markovian Modeling. Although seemingly simple, the memoryless property of the exponential random variable has a crucial technical impact on modeling questions. If one is interested in the evolution of a single-server queue, the number of customers in the queue forms a Markov process only when the arrivals of customers in the queue are Poisson and the service requirements are exponentially distributed. If these assumptions are not met, one can still describe the evolution of the queue as a Markov process if more information is added to the description of the process: if the queue implements the FIFO discipline, it is usually enough to add to the number of customers the residual service time of the customer being served as well as the time till the next arrival; if the queue is Processor-Sharing, then one needs to keep track of the residual service time of each customer, etc...

Although the process giving the number of customers in the queue lives in a countable state space, this is no longer the case when the state descriptor has a continuous component, such as some residual service time. The theory of Markov processes living in an uncountable state space is well developed, but compared to the countable case, the new difficulties introduced are, often, essentially technical. Numerous examples deal with systems with almost arbitrary service distribution, arrival process or even service discipline, but often the high level of technicality introduced obfuscates the main message. In my view it is very valuable to have such examples at hand, since they justify making exponential assumptions, but given a system, one must have a good motivation in order to tackle the problem in all generality. In addition, two other arguments are in favor of the exponential distribution.

First of all, the exponential distribution is met in practice. We have already mentioned the works of Erlang [Erl09, Erl17], where incoming calls are modeled as a Poisson process, and call durations are assumed to follow the exponential distribution. The Poisson approximation essentially comes from the independence of the users who originate the calls, whereas the statistical model for the call durations followed from field measurements. The Poisson approximation turns out to be a very good model for the telephone network, and more generally for most arrival processes which stem from human activity. Sometimes however the Poisson assumption does not apply: In a paper with an enticing title, Paxson and Floyd [PF95] look at the arrival process of packets at a router, and found that it sharply differs from a Poisson process. There is a simple explanation for that, namely that packets in the Internet are generated by computers, so that even if the users' actions are Poisson, each action generates a burst of packets; Paxson and Floyd argue that for different traffic types, the burstiness of the traffic observed cannot be modeled with models derived from Poisson. Another serious problem comes from the long-range dependence of Internet traffic, since long flows induce a correlation of packets over long periods of time. The Poisson process nonetheless remains a realistic model for many arrival processes.

On the other hand it seems harder to justify the exponential assumption when it relates to call durations or more generally to service requirements. Brown *et al.* [**BGM**⁺**05**] for instance investigate in depth the case of a call center, where they conclude that the service times are lognormal. Surprisingly, there exists a remarkable case when this does actually not matter: A network is said to be *insensitive* when the number of customers in steady state at the different nodes depends on the service time distribution only through its mean. For such networks, one can solve the problem for any suitable service distribution, and usually the exponential one is the simplest. A typical example is the aforementioned Erlang B formula: provided calls arrive according to a Poisson process, the number of customers in steady state — and so the blocking probability — only depends on the first moment of the service time distribution, see Takács [**Tak69**] and the references therein. This fundamental property is without any doubt responsible for the enduring success of this simple formula. Insensitivity results were later established for networks of queues by Schassberger [Sch77] and Burman [Bur81], see also the more recent works of Bonald and Proutière [BP02, BP03].

Renormalization Techniques. Even in the "simple" case of the exponential distribution, it is usually challenging to give a simple description of a network's dynamics. Either the network is stable, in which case one wants to know how it stabilizes, or it is unstable, and one wants to know details on the transient paths that lead to infinity. Renormalization techniques aim at providing such a description. The general idea is to take a sequence of suitably renormalized processes that converges: the limiting process then gives insight into the original network's dynamics. We discuss in detail two fundamentals models, the M/M/1 and $M/M/\infty$ queues, which correspond to two different scalings. These two models are central building blocks in the different problems studied therein.

 $\mathbf{M}/\mathbf{M}/\mathbf{1}$ and Fluid Scaling. The M/M/1 queue is the single-server FIFO queue which works at speed one, with Poisson arrivals at rate λ and i.i.d., exponentially distributed with parameter μ , service requirements. To describe this system, one can consider the process (L(t)) that gives the number of customers in the queue; under the above exponential assumptions, this is a Markov process. It lives in the space of non-negative integers $\mathbb{N} = \mathbb{Z}_+ = \{0, 1, \ldots\}$, and its dynamics in the region $\mathbb{N}^* = \{1, 2, \ldots\}$ is simple to describe. On the one hand, it can be seen as the difference of two independent Poisson processes $(\mathcal{N}_{\lambda}(t))$ and $(\mathcal{N}_{\mu}(t))$ with respective parameter λ and μ : if $T_0 = \inf\{t \ge 0 : L(t) = 0\}$ denotes the hitting time of 0, then

(1)
$$L(t) = L(0) + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t), \ 0 \le t \le T_0$$

An equivalent way of describing the dynamics in \mathbb{N}^* is to see (L(t)) as a continuous-time random walk: after a time exponentially distributed with parameter $\lambda + \mu$, the process goes up with probability $\lambda/(\lambda + \mu)$, and otherwise it goes down. As will be discussed below, this analogy is important.

When it hits 0, (L(t)) jumps to 1 after a time exponentially distributed with parameter λ : then using either one of the two descriptions together with the Markov property, it is easy to see that if $\lambda > \mu$, then (L(t)) goes to infinity; if on the other hand $\lambda < \mu$, then it will hit 0 infinitely often. One would like to have a simple description of the dynamics of (L(t)): Equation (1) strongly suggests that starting from $L(0) = \ell$, then $L(t) \approx \ell + (\lambda - \mu)t$, at least for some t > 0. We want to explain how this approximation can be made rigorous. This explanation relies on the estimation of certain important hitting times.

Since the dynamics of (L(t)) is discontinuous at 0, the hitting time T_0 of 0 naturally plays a key role, as Equation (1) highlights. Of course T_0 is random, but if there are many initial customers in the queue, then its behavior can be described with deterministic quantities: if L(0) = n, then as n gets large, T_0 is infinite with high probability if $\lambda > \mu$, whereas T_0/n converges to $t_0 = 1/(\mu - \lambda)$ if $\lambda < \mu$. Hence assuming n large, the approximation $L(t) \approx n + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t)$ holds for all $t \ge 0$ if $\lambda > \mu$, and for $0 \le t \le nt_0$ if $\lambda < \mu$.

Two important remarks should be derived from this simple heuristic. First of all, taking a large initial state makes it possible to give a simpler description of the process; second, in the case $\lambda < \mu$, the relevant time scale is linear in the size of

the initial state. This second point is natural in view of (1): when starting with L(0) = n customers, one will see significant fluctuations in the number of customers when the Poisson processes will have reached values of order of n, which takes a time of order of n as well since a Poisson process grows linearly.

These remarks lead to the so-called *fluid scaling*: consider a sequence $(L_n, n \ge 0)$ of processes such that, for $n \ge 0$, the *n*th process $(L_n(t))$ is an M/M/1 queue starting with *n* customers, i.e., $L_n(0) = n$. Because the time scale is linear in *n*, and because one needs to rescale in *space* by *n* as well in order to avoid a trivial limit as *n* gets large, it is natural to consider the renormalized process $(\overline{L}_n(t))$ defined as follows:

(2)
$$\overline{L}_n(t) = \frac{L_n(nt)}{n}, \ t \ge 0, \ n \ge 1.$$

Because time is sped up by n, the time scale of the renormalized process $(\overline{L}_n(t))$ is the normal time scale, i.e., for this process, one sees significant fluctuations on [0,t] for any t > 0. Because of the above mentioned behavior of T_0 as the size of the initial state gets large, one can prove the following result (see for instance Robert [**Rob03**, Chapter 5] for this and the following results): as n goes to infinity, the sequence of processes $(\overline{L}_n, n \ge 1)$ converges in some sense to the deterministic process, called *fluid limit*, $(x(t) = 1 + (\lambda - \mu)t, t \ge 0)$ when $\lambda > \mu$; when $\lambda < \mu$, then the same convergence holds but only for times $t \le t_0$, i.e., the sequence of processes $(\overline{L}_n, n \ge 1)$ restricted to the time interval $[0, t_0]$ converges to $(x(t), 0 \le t \le t_0)$. This restriction comes from the fact that (1) is only valid for $t \le T_0$, which corresponds to $t \le t_0$ on the fluid scale. An interesting question concerns the behavior of the fluid limit when $t > t_0$ in the case $\lambda < \mu$: once again important hitting times need to be discussed.

Assume $\lambda < \mu$: at time $t_0 = 1/(\mu - \lambda)$, the fluid limit hits 0, i.e., $x(t_0) = 0$. Interpreting the fluid limit as the *n*th system for large *n*, this means, since space has been scaled by a factor *n*, that the number of customers in the *n*th system is much smaller than *n*, it can even be thought of as being 0. Similarly, because of the time and space scaling, the fluid limit will take off from 0 at some time $t > t_0$ only if there are of order of *n* customers in the *n*th system at a time $t > T_0$ still of order of *n*. In other words the hitting time $\tau_n(u) = \inf\{t \ge 0 : L_0(t) \ge un\}$ of level un, for u > 0, starting from 0 (using implicitly the Markov property) is central. If $\tau_n(u)$ is of order *n*, this means that one should see the fluid limit reaching *u* after time t_0 .

Not surprisingly, since the queue is subject to a negative drift when $\lambda < \mu$, $\tau_n(u)$ for u > 0 is much larger than $n: \tau_n(1)$ is of order $(\mu/\lambda)^n \gg n$. Thus the time scale of the fluid scaling does not make it possible to capture such events: on the fluid scale, the process stays stuck at 0. In particular, the sequence of processes (\overline{L}_n) restricted to $[t_0, +\infty)$ converges to the process identically null. Gathering the above observations, one can formulate the fluid approximation in a unified formula for the two cases $\lambda > \mu$ and $\lambda < \mu$:

(3)
$$\overline{L}_n(t) \approx \left(1 + (\lambda - \mu)t\right)^+, \ \forall t \ge 0,$$

where $x^+ = \max(x, 0)$ for any $x \in \mathbb{R}$. This approximation captures the first-order behavior of the process (L(t)): since this comes from the averaging behavior of the large number of initial customers, such approximations are often referred to as functional laws of large numbers. The M/M/1 queue provides a simple and insightful example of such techniques, which are applied in more complex cases, see for instance the book by Bramson [**Bra08**] which studies fluid limits of complex queueing networks; the end goal of this book is the analysis of stability. As will be highlighted later, the fluid scaling plays a special role compared to other scalings due to its importance with respect to stability issues. Other scalings indeed exist: although the fluid scaling of (2) is appropriate for the M/M/1 queue, the right scaling procedure depends in general on the system of interest. For the $M/M/\infty$ queue for instance, the right scaling is Kelly's scaling.

 $\mathbf{M}/\mathbf{M}/\infty$ and Kelly's Scaling. The $M/M/\infty$ is the Markovian queue with an infinite number of servers; in particular, arrivals occur at times of a Poisson process, say with parameter λ , and the process (L(t)) that gives the number of customers is a Markov process. Since there are infinitely many servers, each customer is served immediately upon arrival, and leaves after a time exponentially distributed with parameter μ . The behavior of this queue is radically different from the behavior of the M/M/1: a first and comfortable difference is that the dynamics is continuous at 0. For instance, Kolmogorov's equation has a nice and simple expression,

$$\frac{d}{dt} \left[\mathbb{E}(L(t) \mid L(0) = \ell) \right] = \lambda - \mu \mathbb{E}(L(t) \mid L(0) = \ell),$$

which can be solved to give $\mathbb{E}(L(t) | L(0) = \ell) = \rho + (\ell - \rho)e^{-\mu t}$ with $\rho = \lambda/\mu$.

Besides, the main difference is that the output rate from the queue is not constant nor bounded, since it is proportional to the number of customers in the queue; in particular, and in contrast to the M/M/1, the $M/M/\infty$ cannot be seen as a random walk. Moreover, this new behavior makes that starting from a large initial state L(0) = n, the departures largely outpace the arrivals: in [0, t], the number of arrivals is of order of one (the mean number of arrivals is exactly equal to λt), whereas the number of departures is already of order of n since each one of the n initial customers may have left with positive probability. Thus in contrast to the M/M/1 queue where arrivals and departures occur at the same pace, here there is a different time scale for the arrival and departure processes.

A possibility to get round this problem is to scale the input rate λ by n: then the mean number of arrivals in [0, t] is equal to $n\lambda t$, and is therefore comparable with the number of departures. This simple idea leads to Kelly's scaling, introduced by Kelly [Kel86]. Again we consider a sequence $(L_n, n \ge 1)$ of processes such that the *n*th system represented by $(L_n(t))$ starts with $L_n(0) = n$ customers; in contrast to the fluid scaling however, the input rate λ is scaled by n as well, i.e., the input rate into the *n*th system is equal to $n\lambda$. The rescaled process $(\hat{L}_n(t))$ is obtained by rescaling in space only:

$$\widehat{L}_n(t) = \frac{L_n(t)}{n}, \ t \ge 0, \ n \ge 1.$$

Again, note the difference with the fluid scaling (2) in that time is not sped up. Similarly as for the M/M/1, it can be proved that the sequence of renormalized processes $(\hat{L}_n, n \ge 1)$ converges to the deterministic process $(x(t), t \ge 0)$ which is the unique solution of the following differential equation (see for instance Robert [**Rob03**,

Chapter 6]):

$$\frac{dx(t)}{dt} = \lambda - \mu x(t), \ t \ge 0, \ \text{ and } \ x(0) = 1.$$

This equation is exactly Kolmogorov's equation, which shows that this scaling procedure indeed highlights the first-order behavior of (L(t)). To complete the comparison with the fluid scaling, it is interesting to look at the fluid limit of the $M/M/\infty$.

Imagine one applies the fluid scaling given by (2) to the process (L(t)): note $(\overline{L}_n, n \geq 1)$ the resulting sequence. Similarly as for the M/M/1, the behaviors of T_0 , the hitting time of 0, and of τ_n , the hitting time of n starting from 0, are key to understand the fluid scaling. Because of the exponential decay of (L(t)) reflected by Kolmogorov's equation, it can be proved that T_0 is of order of $\log n$ when L(0) = n gets large. On the other hand the $M/M/\infty$ queue is extremely stable, and not surprisingly, τ_n is very large, it is of order of $(n-1)!(\mu/\lambda)^n$. Since the time scale of the fluid scaling is given by the size of the initial state n, one sees the following: for any t > 0, at times of order nt the process $(L_n(t))$ has already reached 0 and did not have time to bounce back to reach levels of order of n. What this means is that the sequence of processes $(\overline{L}_n, n \geq 1)$ converges on $(0, +\infty)$ to the process identically null. Note that the time t = 0 is excluded from this convergence, since by definition the sequence $(\overline{L}_n(0))$ converges to 1; in particular, the fluid limit of the $M/M/\infty$ exhibits a discontinuity at 0+.

Kelly's scaling can be applied in various settings: Ethier and Kurtz [EK86, Chapter 11] provide some multi-dimensional examples in chemistry and population dynamics — they call these processes *density dependent population processes*. Broadly speaking, Kelly's scaling is suited when events happen at rates which are proportional to the number of customers, so that starting from a large population automatically speeds up time. This is the case for peer-to-peer systems as well, see Massoulié and Vojnović [MV05] or Kesidis *et al.* [KKS09] for instance.

As concluding remark for these two simple yet fundamental examples, note that the fact that the scalings discussed for the M/M/1 and $M/M/\infty$ queues are deterministic and governed by differential equations is quite common. The differential equations usually translate the network's first-order dynamics; the limits are deterministic in "good cases" where the system's main behavior is relatively simple. This is nevertheless not always the case, see for instance Fayolle *et al.* [FIMM91] or Dantzer *et al.* [DHR00] for examples where the limiting process keeps some randomness. In these two examples the process is essentially deterministic, but has from time to time to make a deterministic choice. For instance in [DHR00], the process is piecewise linear, and the slope of each line is randomly chosen when the process hits some boundary.

Fluid Scaling and Stability Analysis. If various scalings can shed light into a system's dynamics, the fluid scaling corresponding to Equation (2) nonetheless plays a special role with respect to stability analysis. In the sequel, the word "stable" refers to a positive recurrent Markov process, when this process lives in a countable state-space; in general, the correct notion of stability is the one of positive Harris recurrence.

Proving the stability of a Markov process usually represents a challenging issue, even for seemingly simple processes and especially for multi-dimensional ones. One approach consists in finding the expression of the stationary distribution: this computational approach is usually intractable, since in most cases there is no closed-form formula for the stationary distribution. A noticeable exception is when the stationary distribution has a product-form, which is true for the large class of reversible or quasi-reversible networks, see the book of Kelly [Kel94] on this subject.

For more complex networks however, it is hopeless to get a hand on the expression of the stationary distribution, and approaches which are more qualitative and less computational are thus needed. Fluid limits provide such an approach, and it is not fortuitous that they were introduced earlier on the simple example of the M/M/1 queue. Indeed, a key feature of the M/M/1 queue, mentioned previously, is that it behaves as a simple random walk in the interior of \mathbb{R}_+ ; its dynamics is nonetheless discontinuous at 0. Many networks can actually be seen this way: for instance, in Jackson's networks the transition from $x \in \mathbb{N}^n$ to $x - e_k + e_\ell$ occurs at rate $\mu_k p_{k,\ell} \mathbb{1}_{\{x_k>0\}}$ for some $\mu_k > 0$, where $p_{k,\ell}$ is the probability for a customer to go from queue k to queue ℓ upon completion of service. It is apparent that this dynamics is that of a random walk in \mathbb{N}^n whose dynamics is characterized by the set of empty queues.

This simple observation motivates Malyshev [Mal93] to study a general class of random walks: the two main conditions satisfied by these random walks is the boundedness of the size of their jump and another *homogeneity* property, namely that the probability to go from one state to another depends only on the distance between these two states and also on the set of empty queues. As any Markov chain, the stability properties of a random walk only depends on what happens far from the origin; moreover, a random walk is a sufficiently nice object so that its long-time behavior is governed by its mean drift. Hence the fluid scaling of Equation (2) is a natural renormalization procedure to apply to a random walk: the *n*th system starts from a large initial state, i.e., far from the origin, and time is sped up in order to bring out the drift. Hence the fluid scaling describes the macroscopic behavior of a random walk started from a large initial state, and so captures the essential features that characterize its stability.

However, not every continuous-time Markov process can be described by means of a random walk: one must be able to go from the continuous time-scale of the Markov process to the discrete one of the random walk, and vice-versa. This essentially depends on the transition rates of the Markov process. When they are bounded, then one can find a universal Poisson process such that transition epochs of the continuous-time Markov process occur at times of this Poisson process: the imbedded Markov chain therefore provides a precise description of the original continuous-time Markov process since they only differ by a time scale which, although random, is smooth. This property is no longer true when transition rates are unbounded: think for instance to the $M/M/\infty$ queue, for which it is not possible to find a Poisson process with a finite intensity that serves as universal clock. This simple remark shows a limitation to this "random walk" approach.

Rybko and Stolyar **[RS92]** develop this approach on an interesting two-node model: they study a multiclass network where they make the connection with fluid limits clearer. Under some static priority rule, their model provides one of the first examples of a stochastic network unstable under the usual conditions $\rho < 1$.
Dai [Dai95] then studied a general class of networks, which can be thought of as Jackson's networks allowing various service disciplines, with several classes of customers, and where service requirements and arrivals follow general distributions. A contribution of this paper is to develop a systematic method to make the link between the stability of the fluid limit and the stability of the original stochastic system, which was already underlying Rybko and Stolyar's analysis. The introduction of this paper provides a good introduction to these questions and to further references. Among others, Dai motivates his work by a previous paper by Dupuis and Williams [DW94] where a reflecting Brownian motion in an orthant is studied. This connection is not surprising in view of the aforementioned results of Malyshev [Mal93]. A recent book by Bramson [Bra08] provides an extensive account on similar models.

In view of the random walk's analogy, a key feature of these models is that the "transition rates" (which make sense under exponential assumptions) are bounded, since transitions of customers within the network are governed by service completion. There are at least two interesting cases which do no fit directly into this framework. A first one concerns the Processor-Sharing discipline, or more generally service disciplines that can serve an unbounded number of customers. There a technical difficulty appears, mentioned by Bramson [Bra08], namely that one needs to control the number of customers finishing their service in any time interval. In such cases it can be convenient to adopt a different technical approach by describing the network with measure valued processes, see for instance Doytchinov *et al.* [DLS01] and Gromoll *et al.* [GPW02] in the case of heavy traffic.

This is nonetheless a rather technical limitation, and another, more natural case is that of systems whose dynamics intrinsically yields unbounded transition rates. This is the case for the $M/M/\infty$, and also for the two models studied in Chapters I and II which indeed inherit salient features of the $M/M/\infty$. In Chapter I the transition rates are unbounded because customers move within a network independently of the service they receive. In this model, motivated by mobile networks, the network's capacity is bounded. In contrast, Chapter II considers a network where the customers act as servers so that, similarly to the $M/M/\infty$, the service capacity is proportional to the number of customers in the network. In this last example, a phenomenon that does not occur in classical queueing networks happens: one needs to deal with time interval of integrable size on which a non-integrable number of events occur.

Since the models studied in Chapters I and II escape the classical framework of queueing networks, it is not possible to directly use fluid limits techniques as developed earlier. In Chapter I these techniques are nonetheless adapted to give stability results on a model for mobile network; the peer-to-peer model of Chapter II is not amenable to such techniques, and Lyapounov type arguments are used. To make these arguments work, a detailed analysis of some *branching processes* is carried out.

Branching Processes. Branching processes are very generic models in applied probability which are used in a wide range of settings: historically they were introduced to study the survival of family names of noble families, but more generally their most natural framework is to represent the evolution of a population.

The Galton-Watson branching process is the simplest branching process: it is a discrete-time Markov process that can be thought of a as a random tree which represents the genealogy of individuals. Imagine a (unisexual) population where independent individuals give birth to a random offspring according to a common offspring distribution. If Z_n is the number of individuals in the *n*th generation for $n \geq 0$, then the sequence (Z_n) satisfies the following recursive equation:

$$Z_{n+1} = \sum_{k=1}^{Z_n} \xi_{k,n}, \ n \ge 0,$$

where $(\xi_{k,n}, k, n \ge 0)$ are i.i.d. random variables distributed according to the common offspring distribution; $\xi_{k,n}$ represents the offspring of the *k*th individual of the *n*th generation. The comprehensive book by Athreya and Ney **[AN72]** presents many results on this process.

A fundamental property of Galton-Watson processes is called the *branching* property: at any time $n \ge 0$, the population starting with x + y individuals can be thought of as coming from two independent branching processes, one starting with x individuals and the other with y individuals. This can be written down as follows: if $(Z_n(x), n \ge 0)$ is a Galton-Watson process started with $Z_0(x) = x$ initial individuals, and $(Z'_n(x), n, x \ge 0)$ follows the same law and is independent of $(Z_n(x), n, x \ge 0)$, then the following distributional equality holds:

$$Z_n(x+y) \stackrel{\text{dist.}}{=} Z_n(x) + Z'_n(y), \quad n \ge 0, \ x, y \in \mathbb{N}.$$

Although seemingly simple, this fundamental property makes it possible to prove many results on Galton-Watson processes; it is even so fundamental that it is the only property (in addition to the strong Markov property) required to define the most general class of branching processes, called *continuous-state branching processes* and introduced by Lamperti [Lam67].

From a modeling standpoint, the Galton-Watson process is limited in that it does not incorporate time: an easy way to add this component is to assign to each edge of the tree representing the branching process i.i.d. labels. Each label then stands for the life of the corresponding individual. In other words, there is a common distribution, say X, such that each individual lives for a duration distributed like X: upon death, the individual splits and gives birth to a random number of new and independent individuals, where the number of new individuals follows the offspring distribution. Such models are called *Bellman-Harris branching processes* following the paper by Bellman and Harris [**BH52**]. They are not so different from Galton-Watson processes in that the time structure is essentially decoupled from the genealogy.

Although more general branching processes exist, such as Crump-Mode-Jagers, Jirina or continuous-state branching processes, Bellman-Harris processes are general enough for our purposes. A *Yule process* is a special kind of Bellman-Harris processes: this is the branching process where particles live for a duration exponentially distributed, and split upon death into exactly two particles; they were initially motivated by the fission of particles.

Due to their intrinsic dynamics, branching processes appear naturally in systems where individuals act identically and independently: this is the case for instance in epidemic models, see for instance Barbour [**Bar09**] where the early stages of an epidemic process is coupled with a branching process. More surprisingly, branching processes appear in other settings as well: in queueing theory they turn out to be useful tools to study the single-server Processor-Sharing queue, see for instance Yashkov [Yas83].

Similarly as epidemic and branching processes are related, Yang and de Veciana **[YdV06]** point out that early stages of a peer-to-peer system under a flash crowd scenario behave similarly as a branching process. This analogy is the starting point of Chapter III, where Yang and de Veciana's ideas are exploited more deeply. Imagine a system that consists, at any time, of a certain number of servers. Each time a customer finishes its service, then it becomes a server in turn, so that there is an increasing and random number of servers; see the previous section presenting peer-to-peer systems for a motivation for such a dynamics. Under a flash crowd scenario, many peers request the file and so servers will not be idle. Under this assumption, servers are independent and each one acts as follows: after a random duration given by the service time of the current customer, it gives birth to a new server — equivalently the server dies and gives birth to two new servers. In other words, the number of servers evolves exactly like a binary Bellman-Harris process, where the life duration of an individual is precisely given by the service time; when this service time is exponentially distributed, then Yule processes naturally appear.

More surprisingly, branching processes turn out to be essential tools in Chapter II where the *stationary* behavior of a peer-to-peer network is studied. In particular, branching processes are used to study a system which, in contrast with epidemic models or peer-to-peer networks under flash crowd, does not exhibit an exponential growth. The analogy is however slightly too complex to be explained here.

We now conclude this section on the mathematical content of this thesis by introducing probabilistic models which, similarly as branching processes, appear naturally in many different problems. Here they appear in connection with the model for a peer-to-peer system under a flash crowd scenario.

Bins and Balls Problems. These problems make it possible to cast many different problems within a somehow unified framework. In the classical version of this problem, n balls are thrown independently into m identical bins: each ball falls in any given urn with probability 1/m. Various asymptotic quantities may be investigated: choosing m as a function of n and letting n go to infinity, one can look at the number of bins which receive at least one ball, or at the number of balls in the bin that receives the most balls; one can wonder how n should be chosen in function of m so that as m goes to infinity, with high probability no bin is empty; etc... These simple models are motivated by classical problems such as the coupon collector's problem or the birthday paradox. The book by Johnson and Kotz [**JK77**] offers a comprehensive account on these questions, see also Chapter 6 of Barbour *et al.* [**BHJ92**] for a recent presentation of these problems.

An extension of these models is when there is an infinite number of bins and a probability vector (p_n) on \mathbb{N} describing the way balls are sent: for $n \ge 0$, p_n is the probability that a ball is sent into the *n*th bin. In one of the first studies in this setting, Karlin **[Kar67]** analyzed the asymptotic behavior of the number of occupied bins. An interesting difference then is that geometrical quantities such as the location of the first empty bin or of the last non-empty one can be investigated. In **[FM85]**, Flajolet and Martin propose a probabilistic algorithm to estimate the cardinality of a multiset: the analysis of this algorithm turns out to be equivalent to estimating the index of the first empty bin when (p_n) is the geometric distribution. The same problem when (p_n) decays as a power law was investigated by Csáki and Földes **[CF76]**, while Hwang and Janson **[HJ08]** look at the number of occupied bins for essentially arbitrary (p_n) .

A further extension of these stochastic models consists in considering random probability vectors (P_n) . These problems have only been investigated recently, and the literature is rather scarce on this topic. A noticeable exception concerns the series of papers of Gnedin *et al.* [Gne04, GINR09, GIR08]: motivated by the problem of integer composition, they analyzed the case where (P_n) decays geometrically fast according to some random variables, i.e., for $n \ge 1$, $P_n = \prod_{i=1}^{n-1} Y_i(1-Y_n)$ where (Y_i) are i.i.d. random variables on (0, 1). Various asymptotic results on the number of occupied bins in this case have been obtained. The random vector (P_n) can be seen as a "random environment" for the bins and balls problem, and it complicates significantly the asymptotic results in some cases. In particular, the indices of the bins in which the balls fall are no longer independent random variables as in the deterministic case, and Chen-Stein's inequality, which makes it possible to tackle many problems in the deterministic setting, does not apply anymore.

Here a bins and balls problem appears in connection with peer-to-peer systems: as explained above, and under certain hypotheses, the population of servers evolves similarly as a Yule process with sequence of split times (t_n) . A sequence (B_i) of i.i.d. exponential random variables describes the times at which customers enter the system, and from a modeling perspective, the first time when two servers are created in a row and no customer arrived in between is important; see Chapter III. Seeing the intervals (t_i, t_{i+1}) as random bins and the points (B_i) as balls, this is exactly a bins and balls problem in random environment. One of the distinctive feature of this work is that using point processes, we are able to describe not only the location of the first empty bin, but the locations of all the first empty bins at once.

Presentation of Chapters

Each chapter of this thesis corresponds to a paper (one of them, corresponding to Chapter II, being at the time of the printing under review):

- **Chapter I:** Florian Simatos and Danielle Tibi. Spatial homogenization in a stochastic network with mobility. To appear in the *Annals of Applied Probability*.
- **Chapter II:** Lasse Leskelä, Philippe Robert, and Florian Simatos. On the stability properties of file-sharing networks. Submitted to *Advances in Applied Probability*, June 2009.
- Chapter III: Florian Simatos, Philippe Robert, and Fabrice Guillemin. A queueing system for modeling a file sharing principle. In *Proceedings of* SIGMETRICS'08, pages 181–192, New York, NY, USA, 2008. ACM.

Chapter IV: Philippe Robert and Florian Simatos. Occupancy schemes associated to Yule processes. Advances in Applied Probability, Vol. 41, Number 2, Pages 600–622, June 2009.

For the sake of completeness, another work that was carried out during my Ph.D. is cited below: the subject is different from the problems treated here, and is therefore not included.

• Florian Simatos. A variant of the Recoil-Growth algorithm to generate multi-polymer systems. *DMTCS Proceedings*, AI:283–294, 2008.

Chapter I: a Model for Mobile Networks. This chapter deals with a stochastic model for mobile networks first investigated by Borst *et al.* [**BPH06**]. The network consists of *n* nodes; $\mu_i \geq 0$ is the service capacity of node *i*, which serves customers according to the Processor-Sharing service discipline. New customers arrive at node *i* according to a Poisson process of intensity λ_i , and then move independently of one another within the network according to a Markov process with *Q*-matrix $Q = (q_{ij})$. In particular, their movements are not governed by the service they receive. Upon arrival, customers generate a service requirement exponentially distributed with mean 1, and are then served according to the Processor-Sharing discipline at each node they visit. Customers leave the network once their service requirement has been fulfilled. The *n*-dimensional process $(X(t) = (X_1(t), \ldots, X_n(t)), t \geq 0)$ where $X_i(t)$ is the number of customers at node *i* at time *t* is then a Markov process whose non-zero transition rates are given, for $i \neq j \in \{1, \ldots, n\}$, by

$$\begin{cases} q_n(x, x + e_i) = \lambda_i \\ q_n(x, x + e_j - e_i) = x_i q_{ij} \\ q_n(x, x - e_i) = \mathbb{1}_{\{x_i > 0\}} \mu_i \end{cases}$$

In these equation e_i denotes the *n*-dimensional vector with every coordinate equal to 0, except the *i*th one equal to 1. The first rate corresponds to an arrival at node *i*; the second one to a transition of a customer from node *i* to node *j*; the last one to a departure from node *i*. The arrival and departure rates are reminiscent of the M/M/1 queue, and the rate for inner transitions of the $M/M/\infty$ queue. This model can be seen as an interacting particles system; with respect to the number of customers, the Processor-Sharing discipline has no impact, but it has nonetheless an appealing queueing motivation. The intuition on this system is the following.

First of all, when no node is empty, the total number of customers locally evolves as an M/M/1 queue with input rate $\lambda = \lambda_1 + \cdots + \lambda_n$ and output rate $\mu = \mu_1 + \cdots + \mu_n$. Thus in view of the fluid approximation (3) of the M/M/1, and as long as there are no empty nodes,

(4)
$$\sum_{i=1}^{n} X_i(t) \approx \left(1 + (\lambda - \mu)t\right)^+.$$

Secondly, imagine that there are many customers in the network. Because the maximal output rate if of order of one — it is bounded by μ — most customers will stay for a long time in the network before leaving. Now assume that Q admits a stationary distribution π , defined by $\pi Q = 0$: since customers stay a long time

in the network, each customer gets close to stationarity and so will after some time be at node i with probability π_i . Since customers' trajectories are moreover independent, the law of large numbers suggests that the fraction of customers at node i at time t is π_i , i.e.,

(5)
$$\frac{X_i(t)}{\sum_{j=1}^n X_j(t)} \approx \pi_i$$

Hence in view of approximations (4) and (5) and of the above discussion, the following approximation is tempting, say when the networks starts off with a large number of customers:

$$X_i(t) \approx \pi_i (1 + (\lambda - \mu)t)^+.$$

Chapter I essentially aims at justifying this approximation, which requires a fine control on the network's behavior. The main idea is that in the fluid regime, customers are instantaneously spread in the network according to π , i.e., approximation (5) indeed holds for t > 0. On the normal time-scale, the time needed to reach this homogenized state is of order of one when there are many customers in the network. Because time is sped up in the fluid regime, this implies that customers are indeed instantaneously homogenized in the fluid regime. On the other hand it is not easy to control that customers stay homogenized for a long time: in the supercritical case $\lambda > \mu$, this is indeed the case thanks to the following almost sure limit that controls the long-term behavior of (X(t)):

$$\lim_{t \to +\infty} \frac{X(t)}{t} = (\lambda - \mu)\pi.$$

In the subcritical case customers stay homogenized as long as there are still many customers in the network, which corresponds to say that customers stay homogenized in the fluid regime as long as fluid limits have not hit 0. After this time, fluid limits stay stuck at 0, which implies stability of the original stochastic system for $\lambda < \mu$.

Although these results seem rather natural, they are actually technical to obtain, and Chapter I is the most technical chapter of this thesis. One of the reasons for that is that we need to control some specific stopping times, typically the time needed by the network to homogenize, starting from an arbitrary state. We control these stopping times thanks to a martingale, whose sole construction is actually one of the achievements of this chapter. Starting from a space-time harmonic function, we show that for some functions F and G, and for any parameter $\alpha > 0$

$$J_{\alpha}(t) = e^{-\alpha t} \int_{\substack{u \in \mathbb{R}^{n-1}: u_i > 0\\ \text{and } \sum_{i=1}^{n-1} u_i < 1}} \prod_{i=1}^n \left(\frac{\tilde{u}_i}{\pi_i}\right)^{X_i(t)} G(\tilde{u}) F(\tilde{u})^{\alpha - 1} du,$$

stopped at some suitable stopping time T_0 is a local martingale, see Theorem 3.1 in Chapter I for the precise notations. The main property of this martingale is that it decouples the time and the state of the system, i.e., it is of the form $e^{-\beta t}V(X(t))$ for some function $V : \mathbb{N}^n \to \mathbb{R}_+$. This form is suitable to yield Laplace transforms of stopping times after using optional stopping arguments. **Chapter II: a Stationary Model for Peer-to-Peer Networks.** In Chapter II, the stability of another model with two different dynamics is investigated. This (n+1)-dimensional model consists of n+1 queues in tandem, labelled from 0 to n, and corresponding to the following non-zero transition rates, for any i = 1, ..., n and any $x = (x_0, ..., x_n) \in \mathbb{N}^{n+1}$:

$$\begin{cases} q(x, x + e_0) = \lambda \\ q(x, x + e_i - e_{i-1}) = \mu_i (x_i \lor 1) \mathbb{1}_{\{x_{i-1} > 0\}} \\ q(x, x - e_n) = \nu x_n. \end{cases}$$

Arrivals occur at rate λ to the first queue, customers stay for an exponential duration in the last queue, and customers in the *i*th queue act as servers for customers in the (i-1)th queue; the boundary condition $x_i \vee 1$ is a technical condition that prevents the existence of absorbing states. This Markov process can model a peerto-peer system for a file with *n* chunks where peers download the chunks sequentially (in order), ask for the next chunk to peers who have one more chunk, and stay for an exponential duration once they have the complete file; this model was motivated by Parvez *et al.* [**PWMC08**].

In the special case n = 2, more general interactions are considered: we find the stability region when the transition rate corresponding to a transition from the first to the second queue is given by

$$q(x, x + e_2 - e_1) = \mu r(x)(x_2 \vee 1) \mathbb{1}_{\{x_1 > 0\}},$$

where $r: \mathbb{N}^2 \to \mathbb{R}_+$ is any function satisfying the condition:

(C)
$$\lim_{x_1 \to +\infty} r(x_1, x_2) = 1, \quad \forall x_2 \ge 0.$$

Considering such functions makes it possible to study variations of the model which offer different insight, for instance:

- **Case** $r(x) = 1 \wedge (x_1/(x_2 \vee 1))$: customers in the second queue act as servers for customers in the first one, but having more servers than customers makes no difference (a customer cannot be served by more than one server).
- **Case** $r(x) = x_1/(x_1 + x_2 \vee 1)$: each customer in the first queue initiates contacts at rate μ , and polls a customer from the first or the second queue uniformly at random. With probability $(x_2 \vee 1)/(x_1 + x_2 \vee 1)$ the customer polled has the file, which the first customer gets instantaneously; downloading times are neglected.

For this system there exists a capacity threshold λ^* : the Markov process is stable if $\lambda < \lambda^*$ and unstable for $\lambda > \lambda^*$. Using Foster's criterion, we prove that λ^* is given by:

$$\lambda^* = \begin{cases} \infty & \text{if } \mu \ge \nu \\ \mu \nu \left[(\nu - \mu)(1 - \log(1 - \mu/\nu)) \right]^{-1} & \text{if } \mu < \nu. \end{cases}$$

In other words, the system can accommodate any input rate if customers share the file long enough (i.e., ν is small); otherwise it can only accommodate a finite input rate.

The technical reason for this dichotomy is the following: broadly speaking, one needs to compare the input into the system to its output, i.e., λ to νx_2 . Thus the

case $x_2 \gg 1$ is a good situation with respect to stability, and so the real bottleneck — when there are many customers — corresponds to $x_1 \gg 1$ and $x_2 \approx 1$. This intuitively explains why the stability region does not depend on the specific function r provided it satisfies Condition (C). If $\mu > \nu$, then the second queue is unstable and grows exponentially fast until it exhausts the first queue; after that, the system essentially resembles an $M/M/\infty$ queue until both queues are almost empty. In the other case $\mu < \nu$, then the second queue is stable, and the mean transition rate from the first to the second queue is then $\mu \mathbb{E}(X_2 \vee 1)$ if $\mathbb{E}(\cdot)$ refers to the expectation with respect to the stationary distribution of X_2 . Hence the system can only accommodate this much throughput.

In the general case $n \ge 3$, it seems difficult to explicitly exhibit the stability threshold. We prove the following result: if

(6)
$$\mu_1 > \mu_2 > \dots > \mu_n - \nu > 0,$$

then the (n+1)-dimensional Markov process is stable for any value of $\lambda > 0$. This case is the equivalent of the good case $\mu > \nu$ in dimension two: informally, customers can be carried over the network very rapidly because all the queues can be build up, thus providing a high throughput. We conjecture that when the above condition fails, then the capacity threshold is finite, and we prove this conjecture in dimension n = 3. In higher dimension similar techniques could work but one needs to face more serious combinatorial problems. Although rather intuitive, the proofs of these results rely on involved couplings with branching processes. Among other results, we need to study a new class of branching processes that we introduce below in order to give the flavor of results encountered in this chapter.

To construct this process, one starts with a Yule process with parameter μ and with a deterministic sequence (σ_n) increasing to infinity. The killed process $(Z(t), t \ge 0)$ is defined by killing a particle of the initial Yule process, and therefore its future progeny with it, at each time σ_n : thus the sequence $(Z(\sigma_n))$ satisfies the following recursion:

$$Z(\sigma_{n+1}) = \left(\sum_{k=1}^{Z(\sigma_n)} \xi_{k,n} - 1\right) \mathbb{1}_{\{Z(\sigma_n) > 0\}}$$

where for each $n \geq 1$, $(\xi_{k,n}, k \geq 1)$ are i.i.d. random variables distributed like $Y(\sigma_{n+1} - \sigma_n)$ where $Y(\cdot)$ is a Yule process with parameter μ and starting with one particle. For $k, n \geq 1$, $\xi_{k,n}$, which is actually a geometric random variable with parameter $e^{-\mu(\sigma_{n+1}-\sigma_n)}$, represents the offspring of the *k*th particle alive at time σ_n in the time interval (σ_n, σ_{n+1}) . Then the killed process $Z(\cdot)$ can survive if and only if

$$\sum_{n\geq 1} e^{-\mu\sigma_n} < +\infty$$

This result is then applied to the model with n queues in tandem to show, in conjunction with Foster's criterion, that it is stable for any input rate $\lambda > 0$ when Condition (6) is satisfied. In dimension n = 3, estimates of the extinction time of this process when Condition (6) fails are performed, which makes it to possible to find the stability region.

Chapters I and II: a Mixture of Two Dynamics. The two models considered in the two first chapters have characteristics both from the M/M/1 and the $M/M/\infty$ queues. Because the transition rates of such Markov processes are unbounded, they do not fit in the usual framework of fluid techniques, see the discussion on fluid limits in the previous section.

A challenging problem in these cases is to decouple the two dynamics: the M/M/1 dynamics is responsible for the stability of the system, while the $M/M/\infty$, which acts on a faster time scale, dictates in some sense how the customers are spread within the network.

Chapter III: a Flash Crowd Scenario. To complement the study of Chapter II on the stationary behavior of file-sharing peer-to-peer networks, we have studied a very simple model for a flash crowd phenomenon. The system starts with one server that offers the file, and N peers that want to download it. Each time a peer finishes downloading the file, it offers it in turn: thus the number of servers increases in time. At time t = 0, every peer initiates an exponential clock with parameter $1/\rho$: when its clock expires, the corresponding peer "wakes up" and enters the system. It is then queued at the server with the smallest number of queued peers, where it is served according to the FIFO discipline and requests a service exponentially distributed. Once served, it becomes a server; thus each time a peer downloads the file, a new server is created where further incoming peers can be queued.

The system starts from a highly overloaded state (N peers want the file, one server offers it); eventually, the situation is reversed, with many servers offering the file and only few peers not awake. The question addressed in Chapter III concerns the time needed for the system to cope with the initial high demand, or equivalently, the time needed to pass from one equilibrium to the other. A first-order approach to this question is the following.

The number P(t) of peers not awake at time t decays exponentially fast, i.e., $P(t) \approx Ne^{-\rho t}$. On the other hand, the number S(t) of servers grows exponentially fast, at least when the system is still overloaded: $S(t) \approx e^t$. The system shifts from one equilibrium to the other when the input rate, proportional to P(t), equals the output rate S(t), i.e., when $S(t) \approx P(t)$, which leads to a time of order $(\log N)/(\rho+1)$. Our conclusion is that this heuristic approach is valid: we nevertheless provide a justification and a more detailed explanation for this approximation. The approach itself is interesting and yields theoretical questions of independent interest, which are treated in Chapter IV.

As a first guess on the time when the system's equilibrium shifts, one can consider the time T_N defined by the first time when two successive servers are created while no peer arrived in between. Intuitively, this amounts to compare the speed at which new peers enter the system to the speed at which new servers are created: T_N corresponds to the first time when servers are created faster than peers arrive.

The time T_N is solely determined by the comparison between two random sequences: note $(A_i, 1 \le i \le N)$ the times at which the N peers wake up (this is a vector of N i.i.d. exponential random variables) and $(\tau_n, n \ge 0)$ the sequence that describes the times at which new servers are created, with $\tau_0 = 0$. Then T_N is defined by

(7)
$$T_N = \inf\{n \ge 0 : A_i \notin (\tau_n, \tau_{n+1}), \ i = 1, \dots, N\}.$$

It can be argued that before the time T_N , empty servers can be safely neglected: a key observation is that if servers were never empty, then their number would evolve exactly like a Yule process, since each server would give birth to a new server after a time exponentially distributed. Hence to get an insight into T_N , one can consider the same problem as (7), but where the sequence (τ_n) is replaced by the sequence (t_n) of split times of a Yule process; for $n \ge 1$, t_n is an accurate approximation of τ_n . One then defines $T_{N,R}$ as the first time that an interval (t_n, t_{n+1}) contains no point (A_i) :

$$T_{N,R} = \inf\{n \ge 0 : A_i \notin (t_n, t_{n+1}), i = 1, \dots, N\}.$$

Then as N gets large, $T_{N,R}$ is a very good approximation of T_N . It turns out that estimating $T_{N,R}$ is not easy — Chapter IV is devoted to solve this problem — and that $T_{N,R}$ is actually not a good indication on the global equilibrium of the network.

Indeed, the first interval (t_n, t_{n+1}) that contains no point does not have a global significance, it is rather due to a rare event, namely that this interval is very small compared to what it ought to be. In terms of peers and servers, this means that the first time that two servers are created in a row and no peer arrived in between is due to the fact that at some point, a server is created very quickly. In particular, around T_N , there are still a great number of peers arriving in between the creation of two successive servers.

Hence T_N does not capture the shift in equilibrium, and this is due to the stochastic fluctuations of the sequence (t_n) . To get rid of this effect, a natural idea is to consider the same problem defining $T_{N,R}$, but instead of looking at the random sequence (t_n) , look at the deterministic sequence $(\mathbb{E}(t_n))$; thus the irrelevant stochastic fluctuations will be avoided. This leads to define the first time $T_{N,D}$ when an interval $(\mathbb{E}(t_n), \mathbb{E}(t_{n+1}))$ contains no point (A_i) :

$$T_{N,D} = \inf\{n \ge 0 : A_i \notin (\mathbb{E}(t_n), \mathbb{E}(t_{n+1})), i = 1, \dots, N\}$$

This problem was already investigated in 1976 by Csáki and Földes **[CF76]**, but we revisited their answer with a modern tool in applied probability, namely Chen-Stein's inequality. This approach gives a more precise answer as well as an indication on the speed of convergence. It is shown that the quantity

$$(\rho+1)T_{N,D} - \log N + \log \log N$$

converges in distribution as N goes to infinity to some non-trivial random variable; in particular $T_{N,D}$ is of order of $(\log N)/(\rho+1)$. Simulations show that this answer is rather good, and gives a good idea of the time when the system's equilibrium shifts. We finally argue, based on simulations, that empty servers cannot be neglected anymore after time $T_{N,D}$. This significantly complicates the analysis since approximating the population of servers with a Yule process is the key idea that allows to derive analytical results.

Chapter IV: Bins and Balls in Random Environment. This chapter is devoted to the study of the bins and balls problem that naturally arises in Chapter III.

A convenient way to describe this bins and balls problem is to consider a random sequence $(t_n, n \ge 0)$, strictly increasing from $t_0 = 0$ to $+\infty$, that divides the positive

half real line \mathbb{R}_+ into random intervals; these intervals play the role of bins. Balls are represented by i.i.d. exponential random variables with parameter $\rho > 0$: the *i*th ball B_i falls in the *n*th interval (t_{n-1}, t_n) if $t_{n-1} < B_i < t_n$.

Because the sequence (t_n) is random, the locations of the balls — i.e., the bins in which they fall — are independent only conditionally on this sequence. Hence conditionally on (t_n) , this is a usual bins and balls problem where the probability P_n for any ball to fall in the *n*th bin is given by

$$P_n = \mathbb{P}(t_{n-1} < B_1 < t_n | (t_k)) = e^{-\rho t_{n-1}} - e^{-\rho t_n} = e^{-\rho t_{n-1}} \left(1 - e^{-\rho (t_n - t_{n-1})}\right).$$

Because the sequence (t_n) is random, so is the probability distribution (P_n) : this is a bins and balls problem in random environment. The random number of balls that fall in the *i*th bin when *n* balls are thrown is denoted $\eta_{i,n}$,

$$\eta_{i,n} = \sum_{j=1}^{n} \mathbb{1}_{\{t_{i-1} < B_j < t_i\}}.$$

Gnedin *et al.* [Gne04] look at the case where (t_n) is a renewal process; due to our motivation coming from Chapter III we investigate the case where the sequence (t_n) is the sequence of split times of a Yule process, i.e.,

$$t_n = \sum_{k=1}^n \frac{E_k}{k}$$

where $(E_k, k \ge 1)$ are i.i.d. exponential random variables with mean 1. With this choice of (t_n) , he random probability distribution (P_n) can be written

$$P_n = \frac{1}{n^{\rho+1}} W_n^{\rho} Z_n,$$

where the sequence (W_n) converges almost surely to some random variable W_{∞} , and (Z_n) converges in distribution to an exponential random variable with parameter ρ . The variable W_{∞} induces a global randomness that affects the whole process, while the variables (Z_n) induce some randomness at the local level only, i.e., at the level of each bin.

Our primary functional of interest is the random point process $\{i \geq 1; \eta_{i,n} = 0\}$ that describes the indices of empty bins: the correct scaling factor — i.e., the order of magnitude of the first empty bin — is $n^{1/(\rho+2)}$, so that we are interested in the convergence of the sequence of point processes (\mathcal{N}_n) where

$$\mathcal{N}_n = \left\{ \frac{i}{n^{1/(\rho+2)}} : i \ge 1, \ \eta_{i,n} = 0 \right\}.$$

It is shown that the sequence (\mathcal{N}_n) converges in distribution as n goes to infinity to the point process

$$\left(W_\infty^{\rho/(\rho+2)}\sigma_i^{1/(\rho+2)}\right)$$

where (σ_i) is a standard Poisson process with parameter $[\rho(\rho+2)]^{-1/(\rho+2)}$. The random variable W_{∞} that induces a global randomness therefore affects the distribution of the limiting process as well. This result follows from the convergence in distribution of the sequence of the two-dimensional point processes (\mathcal{P}_n) where

$$\mathcal{P}_n = \left\{ \left(\frac{i}{n^{1/(\rho+2)}}, nP_i\right) : i \ge 1 \right\}.$$

Let us explain why \mathcal{P}_n is a natural object to study. Conditionally on the environment (t_k) , the average number of balls that fall in the *i*th interval when *n* balls are thrown is given by $\mathbb{E}(\eta_{i,n}|(t_k)) = nP_i$. Thus the first empty bins will be the bins with index *i* such that nP_i is of order 1, which is precisely the indexes on which \mathcal{P}_n sets the focus. Hence \mathcal{P}_n indeed provides a more detailed description of the locations of the first empty bins than \mathcal{N}_n does. Moreover, and perhaps surprisingly, it is technically slightly easier to study the convergence of the sequence (\mathcal{P}_n) than the convergence of (\mathcal{N}_n) .

An interesting phenomenon is studied in conclusion of Chapter IV: due to the non-integrability of $W_{\infty}^{-\rho}$ for $\rho < 1$, the following fact happens for $\rho < 1$. For $\alpha > 0$, define \mathcal{N}_n^{α} by

$$\mathcal{N}_n^{\alpha} = \left\{ \frac{i}{n^{\alpha}} : i \ge 1, \ \eta_{i,n} = 0 \right\}.$$

Informally, \mathcal{N}_n^{α} focuses on empty bins which have indices of order n^{α} . The convergence of the sequence (\mathcal{N}_n) implies that for any x > 0, the sequence $(\mathcal{N}_n^{\alpha}([0, x]))$ converges in distribution to 0 for any $\alpha < 1/(\rho + 2)$. This suggests that the mean number $\mathbb{E}(\mathcal{N}_n^{\alpha}([0, x]))$ of empty bins with index between 0 and xn^{α} should converge to 0 as well, but actually, we have that for $\rho < 1$ and any $1/(2\rho+1) < \alpha < 1/(\rho+2)$:

$$\lim_{n \to +\infty} \mathbb{E} \left(\mathcal{N}_n^{\alpha}([0, x]) \right) = +\infty, \quad \forall x > 0.$$

For $1/(2\rho+1) < \alpha < 1/(\rho+2)$, the discrepancy between the behavior in distribution of the two sequences $(\mathcal{N}_n^{\alpha}([0, x]), n \ge 1)$ and $(\mathbb{E}(\mathcal{N}_n^{\alpha}([0, x])), n \ge 1)$ is due to events which, however of vanishing probability, make the expected value diverge. These rare events are explicitly exhibited and discussed.

Link between Chapters III and IV. These two chapters correspond to two different papers and were written separately, which explains the slight discrepancy in notations between them; for instance, the number of balls is noted N in Chapter III and n in Chapter IV. We hope that this will not hinder their readability. Since these are the most correlated chapters of this thesis, we wish to make clear the connection between them by quickly summarizing their technical content.

Chapter III deals with a bins and balls problem in a deterministic environment: the analysis of the index of the first empty bin relies on Chen-Stein's method, which requires analytic estimates of first and second moments of some random variable. Chapter IV investigates a bins and balls problem in random environment and uses tools from the theory of point processes to describe the locations of the first empty bins. It is mentioned in Chapter IV how the main result obtained in Chapter III can be revisited to give results on point processes, and results from Chapter IV are used in Chapter III to gain some insight into the peer-to-peer model considered.

Organization of References and Citations. We use the following labels. Roman numbers refer to chapters, arabic numbers to statements (lemma, proposition, corollary, theorem, ...) and numbers between parentheses to equations or formulas. For instance, Proposition III.4.1 refers to the proposition with label 4.1 in Chapter III, and Equation (IV.6) to the equation with label (6) in Chapter IV. Since chapters are essentially independent, only few references will be made from one chapter to another and so for simplicity the roman number referring to a chapter is omitted within the same chapter.

At the end of each chapter, the reader will find the list of references used solely in this chapter. A comprehensive list of all the references used in this document can be found on page 163. **References of the Introduction**

- [AFRT06] Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi. Metastability of CDMA cellular systems. In *Mobicom'2006*, Los Angeles, September 2006. ACM.
- [AFRT08] Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi. Stochastic networks with multiple stable points. *Annals of Probability*, 36(1):255–278, 2008.
- [AKM04] Guido Appenzeller, Isaac Keslassy, and Nick McKeown. Sizing router buffers. SIGCOMM Computer Communication Review, 34(4):281–292, 2004.
- [Ald87] D. Aldous. Ultimate instability of exponential back-off protocol for acknowledgment-based transmission control of random access communication channels. *IEEE Transactions on Information Theory*, 33(2):219–223, March 1987.
- [AN72] K. B. Athreya and P. E. Ney. *Branching processes*. Springer, 1972.
- [bAGP+05] Nadia ben Azzouna, Fabrice Guillemin, Stéphanie Poisson, Philippe Robert, Christine Fricker, and Nelson Antunes. Inverting sampled ADSL traffic. In *Proceedings of IEEE ICC'05*, Seoul, 2005.
- [Bar09] A. D. Barbour. Coupling a branching process to an infinite dimensional epidemic process. To appear in ESAIM, 2009.
- [BBP04] T. Bonald, S. C. Borst, and A. Proutière. How Mobility Impacts the Flow-Level Performance of Wireless Data Systems. In *Proceedings of IEEE INFOCOM 2004*, 2004.
- [BFOBR02] N. Benameur, S. Ben Fredj, S. Oueslati-Boulahia, and J. W. Roberts. Quality of service and flow level admission control in the internet. *Computer Networks*, 40(1):57–71, 2002.
- [BFP09] Thomas Bonald, Mathieu Feuillet, and Alexandre Proutière. Is the "Law of the Jungle" sustainable for the Internet? In Proceedings of IEEE INFOCOM 2009, April 2009.
- [BGM⁺05] Lawrence Brown, Noah Gans, Avishai Mandelbaum, Anat Sakov, Haipeng Shen, Sergey Zeltyn, and Linda Zhao. Statistical Analysis fo a Telephone Call Center: A Queueing-Science Perspective. Journal of the American Statistical Association, 100(469):36–50, March 2005.
- [BH52] Richard Bellman and Theodore Harris. On Age-Dependent Binary Branching Processes. *The Annals of Mathematics*, 55(2):280–295, March 1952.
- [BHJ92] A. D. Barbour, Lars Holst, and Svante Janson. Poisson approximation. The Clarendon Press Oxford University Press, New York, 1992.
- [BHP06] A. R. Bharambe, C. Herley, and V. N. Padmanabhan. Analyzing and Improving a BitTorrent Networks Performance Mechanisms. In *Proceedings of IEEE INFOCOM 2006*, pages 2884–2895, Barcelona, Spain, April 2006.
- [Bit] Bittorrent.org, http://bittorrent.org.
- [BMM⁺08] Thomas Bonald, Laurent Massoulié, Fabien Mathieu, Diego Perino, and Andrew Twigg. Epidemic live streaming: optimal performance trade-offs. In *Proceedings of SIGMETRICS'08*, pages 325–336, New York, NY, USA, 2008. ACM.

- [BMP08] Charles Bordenave, David McDonald, and Alexandre Proutière. Performance of random medium access control, an asymptotic approach. In *Proceedings of SIGMETRICS'08*, pages 1–12, New York, NY, USA, 2008. ACM.
- [BP02] Thomas Bonald and Alexandre Proutière. Insensitivity in processorsharing networks. *Performance Evaluation*, 49(1–4):193–209, September 2002.
- [BP03] Thomas Bonald and Alexandre Proutière. Insensitive Bandwidth Sharing in Data Networks. Queueing Systems: Theory and Applications, 44(1):69–100, May 2003.
- [BPH06] Sem Borst, Alexandre Proutière, and Nidhi Hedge. Capacity of wireless data networks with intra- and inter-cell mobility. In *Proceedings* of IEEE INFOCOM 2006, pages 1058–1069, Barcelona, Spain, April 2006.
- [Bra94a] Maury Bramson. Instability of FIFO Queueing Networks. Annals of Applied Probability, 4(2):414–431, 1994.
- [Bra94b] Maury Bramson. Instability of FIFO Queueing Networks with Quick Service Times. Annals of Applied Probability, 4(3):693–718, 1994.
- [Bra08] Maury Bramson. Stability of queueing networks, volume 1950 of Lecture Notes in Mathematics. Springer, Berlin, 2008. Lectures from the 36th Probability Summer School held in Saint-Flour, July 2–15, 2006.
- [Bur81] David Y. Burman. Insensitivity in Queueing Systems. Advances in Applied Probability, 13(4):846–859, 1981.
- [CF76] Endre Csáki and Antónia Földes. On the first empty cell. *Studia Sci*entarum Mathematicarum Hungarica, 11:373–382, 1976.
- [CMP09] D. Chafaï, F. Malrieu, and K. Paroux. On the long time behavior of the TCP window size process. Accepted in Stochastic Processes and their Applications, 2009.
- [Dai95] J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. Annals of Applied Probability, 5(1):49–77, February 1995.
- [DFT06] Olivier Dousse, Massimo Franceschetti, and Patrick Thiran. On the throughput scaling of wireless relay networks. *IEEE Transactions on Information Theory (joint issue with IEEE/ACM Transactions on Networking)*, 52(6):2756–2761, 2006.
- [DGR02] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. A Markovian analysis of Additive-Increase Multiplicative-Decrease (AIMD) algorithms. Advances in Applied Probability, 34(1):85–111, 2002.
- [DHR00] Jean-François Dantzer, Mostapha Haddani, and Philippe Robert. On the stability of a bandwidth packing algorithm. *Probability in the En*gineering and Informational Sciences, 14(1):57–79, 2000.
- [DLS01] Bogdan Doytchinov, John Lehoczky, and Steven Shreve. Real-time queues in heavy traffic with earliest-deadline-first queue discipline. Annals of Applied Probability, 11(2):332–378, 2001.
- [Dum97] Vincent Dumas. A Multiclass Network with Non-Linear, Non-Convex, Non-Monotonic Stability Conditions. *Queueing Systems: Theory and Applications*, 25(1–4):1–43, 1997.

[DW94]	P. Dupuis and R. J. Williams. Lyapounov functions for semimartingale reflecting Brownian motions. <i>Annals of Applied Probability</i> , 22(2):680–702, 1994
[EK86]	Stewart N. Ethier and Thomas G. Kurtz. <i>Markov Processes, Charac-</i> <i>terization and Convergence</i> . Wiley Series in Probability and Mathe- matical Statistics, 1986.
[Erl09]	A. K. Erlang. Sandsynlighedsregning og Telefonsamtaler. Nyt Tidsskrift for Matematik B, 20:33, 1909.
[Erl17]	A. K. Erlang. Løsning af nogle Problemer fra Sandsynlighedsregningen af Betydning for de automatiske Telefoncentraler. <i>Elektroteknikeren</i> , 13:5, 1917.
[FDTT07]	M. Franceschetti, O. Dousse, D. N. C. Tse, and P. Thiran. Closing the gap in the capacity of random wireless networks via percolation theory. <i>IEEE Transactions on Information Theory</i> , 53(3):1009–10018, 2007.
[FIMM91]	G. Fayolle, I.A. Ignatyuk, V.A. Malyshev, and M.V. Men'shikov. Ran- dom walks in two-dimensional complexes. <i>Queueing Systems, Theory</i> and Applications, 9(3):269–300, 1991.
[FJ93]	Sally Floyd and Van Jacobson. Random early detection gateways for congestion avoidance. <i>IEEE/ACM Transactions on Networking</i> , 1(4):397–413, 1993.
[FM85]	Philippe Flajolet and G. Nigel Martin. Probabilistic counting algorithms for data base applications. <i>Journal of Computer and System Sciences</i> , 21:182–209, 1985.
[GINR09]	Alexander V. Gnedin, Alex Iksanov, Pavlo Negadajlov, and Uwe Roesler. The Bernoulli Sieve revisited. To Appear in the <i>Annals of</i> <i>Applied Probability</i> , 2009.
[GIR08]	Alexander V. Gnedin, Alex Iksanov, and Uwe Roesler. Small parts in the Bernoulli sieve. <i>DMTCS proceedings</i> , AI:235–242, 2008.
[GK00]	P. Gupta and P. R. Kumar. The capacity of wireless networks. <i>IEEE Transactions on Information Theory</i> , 46(2):388–404, March 2000.
[GMR06]	Christos Gkantsidis, John Miller, and Pablo Rodriguez. Anatomy of a P2P Content Distribution system with Network Coding. In <i>Proceed-</i> <i>ings of the 5th International Workshop on Peer-to-Peer Systems</i> , Santa Barbara, CA, USA, February 27–28 2006.
[Gne04]	Alexander V. Gnedin. The Bernoulli Sieve. <i>Bernoulli</i> , 10(1):79–96, 2004.
[GPW02]	H. Christian Gromoll, Amber L. Puha, and Ruth J. Williams. The fluid limit of a heavily loaded processor sharing queue. <i>Annals of Applied Probability</i> , 12(3):797–859, 2002.
[Gri99] [GT01]	Geoffrey Grimmett. <i>Percolation</i> . Springer, 1999. M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In <i>Proceedings of IEEE INFOCOM 2001</i> , volume 3, pages 1360–1369, 2001.
[GV04]	A. T. Gai and L. Viennot. Broose: A practical distributed hashtable based on the de-bruijn topology. In <i>International Conference on Peerto-Peer Computing (P2P)</i> , Zurich, August 2004.

- [HJ08] Hsien-Kuei Hwang and Svante Janson. Local limit theorems for finite and infinite urn models. *Annals of Probability*, 36(3):992–1022, 2008.
- [JK77] Norman L. Johnson and Samuel Kotz. Urn models and their application: an approach to modern discrete probability theory. New York: Wiley, 1977.
- [Kar67] Samuel Karlin. Central limit theorems for certain infinite urn schemes. Journal of Mathematics and Mechanics, 17:373–401, 1967.
- [Kel86] F. P. Kelly. Blocking probabilities in large circuit-switched networks. Advances in Applied Probability, 18(2):473–505, June 1986.
- [Kel94] F. P. Kelly. Reversibility and Stochastic Networks. Wiley, Chichester, 1979, reprinted 1987, 1994.
- [KH95] Raymond Knopp and Pierre A. Humblet. Information capacity and power control in single-cell multiuser communications. In ICC'95, IEEE International Conference on Communications, June 18–22, 1995 – Seattle, USA, June 1995.
- [KKS09] G. Kesidis, T. Konstantopoulos, and P. Sousi. A stochastic epidemiological model and a deterministic limit for BitTorrent-like peer-to-peer file-sharing networks. Springer Lecture Notes in Computer Science, 5425:26–36, 2009.
- [Lam67] John Lamperti. The Limit of a Sequence of Branching Processes. *Probability Theory and Related Fields*, 7(4):271–288, August 1967.
- [Loy62] R. M. Loynes. The stability of a queue with non-independent interarrival and service times. Mathematical Proceedings of the Cambridge Philosophical Society, 58(03):497–520, 1962.
- [Mal93] V. A. Malyshev. Networks and dynamical systems. Advances in Applied Probability, 25(1):140–175, 1993.
- [MR99] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. In *Proceedings of IEEE INFOCOM 1999*, volume 3, pages 1395–1403 vol.3, March 1999.
- [MT08] Laurent Massoulié and Andrew Twigg. Rate-optimal schemes for Peerto-Peer live streaming. *Performance Evaluations*, 65(11-12):804–822, 2008.
- [MV05] Laurent Massoulié and Milan Vojnović. Coupon replication systems. In Proceedings of SIGMETRICS'05, volume 33, pages 2–13, Banff, Alberta, Canada, June 2005.
- [PF95] Vern Paxson and Sally Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3:226–244, 1995.
- [PGES05] Johan Pouwelse, Pawel Garbacki, Dick Epema, and Henk Sips. The Bittorrent P2P File-sharing System: Measurements and Analysis. In Proceedings of IPTPS'05, Ithaca, NY, February 2005.
- [PWMC08] Nadim Parvez, Carey Williamson, Anirban Mahanti, and Niklas Carlsson. Analysis of bittorrent-like protocols for on-demand stored media streaming. In *Proceedings of SIGMETRICS'08*, pages 301–312, New York, NY, USA, 2008. ACM.
- [QS04] Dongyu Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proceedings of SIGCOMM'04*, pages 367–378, New York, NY, USA, August 2004. ACM.

[Rob03]	Philippe Robert. <i>Stochastic Networks and Queues.</i> Stochastic Modelling and Applied Probability Series. Springer-Verlag, New York, 2003. xvii+398 pp.
[RS92]	A. N. Rybko and A. L. Stolyar. On the ergodicity of random processes that describe the functioning of open queueing networks. <i>Problems on Information Transmission</i> , 28(3):3–26, 1992.
[SAV06]	Riikka Susitaival, Samuli Aalto, and Jorma Virtamo. Analyzing the dynamics and resource usage of P2P file sharing by a spatio-temporal model. <i>International Workshop on P2P for High Performance Comp</i> - tutational Sciences (P2P-HPCS'06) in conjunction with ICCS, 2006.
[Sch77]	R. Schassberger. Insensitivity of Steady-State Distributions of Gen- eralized Semi-Markov Processes. Part I. <i>The Annals of Probability</i> , 5(1):87–99, 1977.
[SG07]	Walid Saddi and Fabrice Guillemin. Measurement based modeling of eDonkey peer-to-peer file sharing system. In <i>Proceedings of the Inter-</i> <i>national Teletraffic Congress 20</i> . Ottawa, June 2007.
[Tak69]	Lajos Takács. On Erlang's formula. Annals of Mathematical Statistics, 40(1):71–78, 1969.
[TDG08]	Dominique Tschopp, Suhas Diggavi, and Matthias Grossglauser. Hier- archical routing over dynamic wireless networks. <i>SIGMETRICS Per-</i> <i>formance Evaluation Review</i> , 36(1):73–84, 2008.
[TWN06]	Ye Tian, Di Wu, and Kam Wing Ng. Modeling, Analysis and Improvement for BitTorrent-Like File Sharing Networks. In <i>Proceedings</i> of <i>IEEE INFOCOM 2006</i> , pages 2921–2931, Barcelona, Spain, April 2006.
[Wal09]	N. S. Walton. Proportional Fairness and its Relationship with Multi- class Queueing Networks. Preprint available on arXiv, 2009.
[Yas83]	S. F. Yashkov. A Derivation of Response Time Distribution for an M/G/1 Processor-Sharing Queue. Problems of Control and Information Theory, 12(2):133-148, 1983.
[YdV06]	Xiangying Yang and Gustavo de Veciana. Performance of peer-to-peer networks: service capacity and role of resource sharing policies. <i>Per-</i> <i>formance Evaluations</i> , 63(3):175–194, 2006.
[ZJM00]	Hui Zang, Jason P. Jue, and Biswanath Mukherjee. A review of rout- ing and wavelength assignment approaches for wavelength-routed opti- cal WDM networks. <i>Optical Networks Magazine</i> , 1(1):47–60, January 2000.

CHAPTER I

Spatial Homogenization in a Stochastic Network with Mobility

1.	Introduction	43	
2.	Framework and Notations	46	
3.	Martingale	49	
4.	Two Key Representations	52	
5.	The Space Renormalized Process	54	
6.	The Supercritical Regime	58	
7.	Stability of the Subcritical Regime	66	
Appendix A. Martingale Construction			
References of Chapter I			

1. Introduction

Recent wireless technologies have triggered interest in a new class of stochastic networks, called *mobile networks* in the technical literature [BPH06, GT01]. In contrast with Jackson networks where users move upon completion of service at some node, in these mobile networks, transitions of customers within the network occur *independently* of the service received. Moreover, at any given time, each node capacity is divided between the users present, whose service rate thus depends on the capacity and on the state of occupancy of the node. Once his initial service requirement has been fulfilled, a customer definitively leaves the network. In Borst et al. [BPH06], complex capacity sharing policies are considered, but in the simplest setting, which will be of interest to us, nodes implement the Processor-Sharing discipline by dividing their capacity equally between all the users present. Previous works [BPH06, GT01] have mainly focused on determining the stability region of such networks, and it has been commonly observed that the users' mobility represents an opportunity for the network to increase this region. Indeed, because of their mobility, users offer a diversity of channel conditions to the base stations (in charge of allocating the resources of the nodes), thus allowing them to select the users in the most favorable state. Such a scheduling strategy is sometimes referred to as an opportunistic scheduling strategy, see Borst [Bor05] and the references therein for more details.

In the present chapter, we investigate from a mathematical standpoint a basic Markovian model for a mobile network, derived from Borst *et al.* [BPH06]. In this simple setting, customers arrive in the network according to a Poisson process with intensity λ , and move independently within the network, according to some Markovian dynamics with a common rate matrix Q. Service requirements are exponentially distributed with mean 1, and customers are served at each node they visit according to the Processor-Sharing discipline, until their demand has been satisfied. The total capacity of the network, defined as the sum of all the individual capacities of the nodes, is denoted by μ . It corresponds to the instantaneous output rate of the network when no node is empty, i.e., when there is at least one customer at each node.

It is of particular interest to note that, even if Q is reversible, because of the arrival and departure processes, the system is not reversible. This contrasts with earlier works in which particle systems with similar dynamics have been investigated under reversibility assumptions. In Caputo and Posta [**CP07**], the authors look at a closed system (i.e., with parameters $\lambda = \mu = 0$) where transition rates are chosen such as to yield a reversible dynamics. In this case, the stationary distribution of the system has a product form, and the authors are interested in showing that the convergence to equilibrium is exponentially fast. Their approach essentially relies on logarithmic Sobolev type inequalities.

In our case however, a different set of questions is addressed, involving different tools. Since the system under consideration is open, it may be unstable, so that a natural issue is to determine the stability region. We prove, as was conjectured in Borst *et al.* [BPH06], that the intuitive, simple condition $\lambda < \mu$ is indeed the stability condition (the critical case $\lambda = \mu$ is not considered). In contrast with Jackson networks for which the stability condition is local, in the sense that each node has to satisfy some constraint, here only the global quantities λ and μ matter. This shows that mobility allows to make the most of the potential service capacity of the network, corroborating the results previously mentioned. Note that $\lambda < \mu$ being a necessary condition is obvious, since μ is the maximal output rate. But surprisingly, proving that it is sufficient requires very technical tools, among which the use of fluid limits and martingale techniques. In particular, the long and tedious Appendix A of this chapter is solely devoted to the construction of a martingale which provides key estimates for showing that $\lambda < \mu$ corresponds to a stable system.

This martingale is a multidimensional (therefore complicated) generalization of the martingale built in Fricker *et al.* [FRT99] for the $M/M/\infty$ queue, and this is not completely surprising, since as will be seen, the model inherits salient properties of the $M/M/\infty$ queue. Besides, the construction of a martingale associated to a multidimensional process represents one of the technical achievements of this chapter: such examples are indeed pretty scarce in the literature. Similarly as in Fricker *et al.* [FRT99], the approach relies on building a family of space-time harmonic functions indexed by some parameter $c \in \mathbb{R}^n$, and then on integrating over cin such a way as to preserve the harmonic property.

Through studying both the stability region and the unstable regime, a detailed description of the behavior of the system is given, resulting in two versions (stable and unstable) of the following rough property: When many users are present in the network, they get approximately distributed among the nodes according to the unique invariant distribution π associated to Q, the latter being assumed

irreducible. It must be emphasized that yet, contrary to Borst *et al.* [BPH06], customers' movements are not assumed stationary.

As a first argument for this spatial homogenization, the law of large numbers suggests that, when the total number of users initially present in the network is large, the proportions of users at the different nodes should be close to π after some time, related to the convergence to π of the Markov process associated to Q. The more delicate question, that next arises, of how long these proportions stay close to π constitutes the main challenging issue of the chapter, that requires martingale techniques for estimating the deviation time from π .

The short term reach of π is understandable from an analogy with the $M/M/\infty$ queue: indeed, independence of the customers' trajectories yields that, similarly to the $M/M/\infty$ queue, the output rate from any node due to inner transitions is directly proportional, through Q, to the number of customers at this node. When the network is overloaded, the relative occupancies of the nodes should then, after a while, be close to the internal traffic balance ratios, given by π .

A more explicit analogy with another classical queueing model is provided by the following simple but crucial observation: As long as no node is empty, the total number of customers simply evolves as an M/M/1 queue with input rate λ and output rate μ . And this is in particular the case when the distribution of customers is close to π . This interplay between, on the one hand, the proportions of customers at the different nodes, and on the other hand their total number, will underly the analysis all along the chapter.

While the short term behavior, which results in the spreading of customers according to π , is dominated by the $M/M/\infty$ dynamics, the long term behavior is essentially driven by the M/M/1 dynamics of the total number of customers. This naturally suggests that two different scalings have to be considered: one, corresponding to the $M/M/\infty$ dynamics, where only space is scaled, and not time; and a second one, where both space and time are scaled, corresponding to the fluid scaling of the M/M/1 queue. Note that the natural scaling for the $M/M/\infty$ queue is the so-called Kelly scaling, in which space and input rate are scaled. Here, since the input rate at each node due to inner transitions is a linear function of the numbers of customers at the different nodes, there is no need to scale the external input rate λ . Inner movements dominate the dynamics and the space scaled process converges, analogously to the $M/M/\infty$ queue under Kelly's scaling, to some deterministic trajectory, with limit point at infinity here given by π .

The coexistence of these two different scalings makes the use of fluid limits both original and challenging. Fluid limits are a standard tool in the analysis of complicated stochastic networks. Rybko and Stolyar [**RS92**] is one of the first papers using this technique together with Dai [**Dai95**]. Dupuis and Williams [**DW94**] presented similar ideas in the context of diffusions. In a series of papers Bramson [**Bra96a**, **Bra96b**] describes the precise evolution of fluid limits for various queueing networks. See also the books by Chen and Yao [**CY01**] and Robert [**Rob03**]. In the context of networks, fluid limits have been used mainly for Markov processes which behave locally as random walks. For this reason, results related to fluid limits are sometimes presented as functional laws of large numbers. Because of the mixture of two different dynamics, given by the M/M/1and $M/M/\infty$ models, our framework is somewhat different. A second important difference with the existing literature concerns tightness results which are usually easy to obtain, mainly because transition rates are generally bounded: this not the case here.

The long term analysis is twofold. Deriving fluid limits requires a control on the process over time periods of the same order as the initial number of customers (since the fluid scaling parameter is the same for time and space). In the stable case this is obtained by showing that the deviation time from π is essentially larger than the time for the underlying M/M/1 queue to empty. The unstable case exhibits a more striking behavior: the deviation time from π is not only large compared to the initial number of customers, but is even *infinite* with high probability. This amounts to a control of the whole trajectory: the distribution of users among nodes stays trapped in any neighborhood of π with high probability as the initial state is large. This result is related to a strong convergence result stating that, for any fixed (non scaled) initial state, the system almost surely diverges along the direction of π . A similar phenomenon has been exhibited in Athreya and Kang [**AK98**], in the context of branching Markov chains, i.e., Galton-Watson branching processes where individuals located at some countable set of sites move at their birth time.

These various remarks and outline of results lead to the following organization for the chapter. Section 2 gives a precise description of the stochastic model and introduces the notations that will hold throughout the chapter. We have already mentioned the construction of a martingale which gives important estimates through optional stopping techniques: Section 3 introduces this martingale, and provides the main estimate that will be used. Due to its technicality, the construction of the martingale is postponed to the Appendix A.

Section 4 establishes a decomposition of the process as, mainly, the difference between two processes of the same type but with no departures. For such a process (with null service capacity), a representation involving labelled particles is given. Both representations will help derive the almost sure convergence result of Section 6.

The three last sections are devoted to analyzing the behavior of the system. Section 5 deals with the short term behavior, thus studying the only space renormalized process. Section 6 studies the supercritical case $\lambda > \mu$, establishing among other results the almost sure convergence of the proportions to the equilibrium distribution π as $t \to \infty$. Finally, Section 7 proves the stability of the system in the subcritical case $\lambda < \mu$.

2. Framework and Notations

This section gives a precise description of the model under consideration and introduces the main notations. The network is described by a Markov process $X = (X(t), t \ge 0)$ characterized by its infinitesimal generator, given by (1) below.

Section 6 will make use, in the particular case of null service capacity, of a more explicit representation of X involving a sequence of Markov jump processes that represent the trajectories of the successive customers entering the network. The general description of the system through its Markovian dynamics provided in the present section is however sufficient for most results of the chapter, especially for

building a family of martingales and for determining the stability condition.

The network consists of n nodes between which customers perform independent (continuous time) Markovian routes during their service. In this setting, transitions of customers from one node to another are driven by some rate matrix $Q = (q_{ij}, 1 \le i, j \le n)$ and are thus not triggered by service completion.

New customers arrive at node i = 1, ..., n according to a Poisson process with intensity $\lambda_i \geq 0$, and then move independently according to the Markovian dynamics defined by Q. The arrival processes at the different nodes are independent, so that the global arrival process is Poisson with intensity $\lambda = \sum_{i=1}^{n} \lambda_i$. The case $\lambda = 0$ corresponds to a system with only initial customers, and no new arrivals.

Upon arrival, or at time t = 0 for those initially present, customers generate a service requirement which is exponentially distributed with mean 1. All service requirements, arrival processes and Markovian routes are assumed to be mutually independent.

Node $i, 1 \leq i \leq n$, has service capacity $\mu_i \geq 0$, which is divided at any time between the customers present, according to the Processor-Sharing discipline: If Nis the number of customers present at node i, then each of these N customers is served at rate μ_i/N . The service rate of a given customer thus evolves in time, depending on his current position and on its occupancy level. Once a customer has received a service that meets his initial requirement, he leaves the network.

The total service capacity of the network is defined as $\mu = \sum_{i=1}^{n} \mu_i$. Notice that, due to the exponential nature of the services, the mechanism of departure from one node by completion of service does not distinguish the present Processor-Sharing discipline from the FIFO discipline: the instantaneous output rate from the system at node *i* is μ_i provided that node *i* is not vacant. The total output rate is then μ when no node is empty.

The process of interest is $X = (X(t), t \ge 0)$ defined by

$$X(t) = (X_1(t), \dots, X_n(t)), \quad t \ge 0,$$

where $X_i(t)$, for i = 1, ..., n, is the number of customers present at node i at time t. The Markovian nature of the movements together with the exponential assumption for the service distribution imply that X is a Markov process in \mathbb{N}^n with infinitesimal generator Ω given, for any function $f: \mathbb{N}^n \to \mathbb{R}$ and any $x = (x_1, \ldots, x_n) \in \mathbb{N}^n$, by

(1)
$$\Omega(f)(x) = \sum_{i=1}^{n} \lambda_i \big(f(x+e_i) - f(x) \big) + \sum_{i=1}^{n} \mathbb{1}_{\{x_i > 0\}} \mu_i \big(f(x-e_i) - f(x) \big) \\ + \sum_{1 \le i \ne j \le n} q_{ij} x_i \big(f(x+e_j - e_i) - f(x) \big),$$

where $e_i \in \mathbb{N}^n$ has all coordinates equal to 0, except for the *i*th one, equal to 1.

The introduction has highlighted that this system is a mixture of two classical models in queueing theory, the M/M/1 and the $M/M/\infty$ queues. This is readable in the expression of the generator given in (1), where the two first sums are reminiscent of the M/M/1 queue, and the last one of the $M/M/\infty$ queue.

The rate matrix Q is assumed to be irreducible, admitting $\pi = (\pi_i, 1 \le i \le n)$ as its unique stationary distribution, characterized by the relation For technical reasons related to the construction of the martingale introduced in Section 3 (see the Appendix A), we require the additional assumption that Q is diagonalizable. This assumption is satisfied if Q is reversible with respect to π , but it is in general a much less restrictive constraint.

For any $t \ge 0$, the random vector X(t) will often be described in terms of the total number of customers L(t) and the proportions of customers at the different nodes $\chi(t) = (\chi_i(t), 1 \le i \le n)$. More formally, define

$$L(t) = \sum_{j=1}^{n} X_j(t) = |X(t)| \text{ and } \chi_i(t) = \frac{X_i(t)}{L(t)}, \quad 1 \le i \le n, \ t \ge 0,$$

with the convention that $\chi(t) = e_1$ when L(t) = 0. Here, and more generally for any $x = (x_1, \ldots, x_n) \in \mathbb{R}^n$, |x| denotes the ℓ^1 norm in \mathbb{R}^n : $|x| = \sum_{i=1}^n |x_i|$.

The vector $\chi(t)$ can be identified with a probability measure on $\{1, \ldots, n\}$: namely, the empirical distribution of the positions of the L(t) customers present in the network at time t. Denote by

$$\mathcal{P} = \left\{ \rho \in [0, +\infty[^n : \sum_{i=1}^n \rho_i = 1] \right\}$$

the state space of $\chi(t)$. The interior set of \mathcal{P} is $\mathring{\mathcal{P}} = \{\rho \in]0, +\infty[^n: \sum_{i=1}^n \rho_i = 1\}.$

As emphasized earlier, the deviation of $\chi(t)$ from π will be of particular interest in the forthcoming analysis. It will be measured, depending on circumstances, by the ℓ^{∞} distance $\|\chi(t) - \pi\|$:

$$\|x\| = \max_{1 \le i \le n} |x_i|, \quad x = (x_1, \dots, x_n) \in \mathbb{R}^n,$$

or by the relative entropy $H(\chi(t), \pi)$, where $H(\cdot, \pi)$ is defined on the set \mathcal{P} of probability measures on $\{1, \ldots, n\}$, by

$$H(\rho, \pi) = \sum_{i=1}^{n} \rho_i \log \frac{\rho_i}{\pi_i} \in [0, +\infty[, \rho \in \mathcal{P}.$$

For $t \ge 0$, the quantity $H(\chi(t), \pi)$ will also be more simply denoted H(t). The process $(H(t), t \ge 0)$ will spontaneously appear in the expression of the key martingale J_{α} introduced in the next section.

The different deviation times of $\chi(t)$ from π , or conversely, the time needed for $\chi(t)$ to reach a given neighborhood of π , will be of particular interest. For any $\varepsilon > 0$, T_{ε} (resp. T^{ε}) denotes the first time when the ℓ^{∞} distance between $\chi(t)$ and π is smaller (resp. larger) than ε :

$$T_{\varepsilon} = \inf\{t \ge 0 : \|\chi(t) - \pi\| \le \varepsilon\} \text{ and } T^{\varepsilon} = \inf\{t \ge 0 : \|\chi(t) - \pi\| > \varepsilon\}.$$

Most results will be written down in terms of these two stopping times, but it will be sometimes more convenient to work with the deviation time T_H^{ε} from π in terms of the relative entropy:

$$T_H^{\varepsilon} = \inf\{t \ge 0 : H(t) > \varepsilon\}.$$

All results on deviation times of $\chi(t)$ from π defined in terms of the ℓ^{∞} distance $\|\chi(t) - \pi\|$ can be translated into analogous estimates in terms of the relative entropy H(t) thanks to the following classical result:

LEMMA 2.1. There exist two π -depending positive constants C_1 and C_2 such that, for all $\rho \in \mathcal{P}$:

$$C_1 \|\rho - \pi\|^2 \le H(\rho, \pi) \le C_2 \|\rho - \pi\|^2.$$

In particular, for any $\varepsilon > 0$, $T_H^{C_1 \varepsilon^2} \le T^{\varepsilon} \le T_H^{C_2 \varepsilon^2}.$

Another stopping time will play a central role: namely the first time, denoted by \mathcal{T}_0 , when the system has an empty node. Formally,

 $\mathcal{T}_0 = \inf\{t \ge 0 : \exists i \in \{1, \dots, n\}, \ X_i(t) = 0\}.$

Indeed, the martingale property for the family of integrals presented in Section 3 will hold only up to time \mathcal{T}_0 , i.e., as long as the output rate at each node *i* is exactly equal to μ_i . In the same way, it will be easily shown that, for $t < \mathcal{T}_0$, L(t) behaves exactly like the M/M/1 queue with input rate λ and output rate μ .

A last useful remark concerning these stopping times is that, when \mathcal{T}_0 is finite, $\|\chi(\mathcal{T}_0) - \pi\| \ge \min \pi_i (> 0)$. Together with Lemma 2.1, this immediately gives the following result:

LEMMA 2.2. There exists $\varepsilon_0 > 0$ such that $T^{\varepsilon} \vee T_H^{\varepsilon} \leq \mathcal{T}_0$ holds for any $\varepsilon \leq \varepsilon_0$.

3. Martingale

The results of this section are twofold: Theorem 3.1 gives the (almost) explicit expression of a local martingale $J_{\alpha}(\cdot \wedge \mathcal{T}_0)$, indexed by some positive parameter α , and Proposition 3.2 derives the main estimate on deviation times T_H^{ε} of $\chi(t)$ from π , that will be used in Sections 6 and 7. Concerning the construction of J_{α} , the present section only aims at giving the main lines. The (numerous) technical details are postponed to the Appendix A.

The approach for constructing the martingale J_{α} is similar to the approach used in Fricker [**FRT99**] for the $M/M/\infty$ queue. The idea is to first exhibit a family of space-time harmonic functions $(h_v(t, x), v \in \mathbb{R}^n)$ for the generator Ω given by (1), and then to integrate $h_v(t, x)f(v)$ with respect to v for some suitable function f, on some well chosen time dependent domain. The last step is then to make a change of variables so that the new harmonic function is split into two factors, respectively depending on time and space. The resulting local martingale is then adapted for an optional stopping use, leading to hitting times estimations.

Some notations are required at this point. Denote by $(P_t, t \in \mathbb{R})$ the *Q*-generated Markov semi-group of linear operators in \mathbb{R}^n : $P_t = e^{tQ}$, extended to all real indices t into a group. For $v \in \mathbb{R}^n$ and $t \in \mathbb{R}$, define

$$\phi(v,t) = (\phi_i(v,t), 1 \le i \le n) = P_{-t}v.$$

Theorem 3.1 below requires the technical assumption that Q is diagonalizable. Let θ be the trace of -Q, so that $\theta > 0$, and let $S \subset \mathbb{R}^{n-1}$ be the projection of $\mathring{\mathcal{P}} \subset \mathbb{R}^n$ on the n-1 first coordinates, i.e.,

$$\mathcal{S} = \left\{ u = (u_1, \dots, u_{n-1}) \in \mathbb{R}^{n-1} : \forall i = 1, \dots, n-1, u_i > 0 \text{ and } \sum_{i=1}^{n-1} u_i < 1 \right\}.$$

For any $u \in S$, denote by $\tilde{u} \in \mathring{\mathcal{P}}$ the *n*th dimensional vector which completes *u* into a probability distribution, i.e., $\tilde{u}_i = u_i$ for any $1 \le i \le n-1$ and $\tilde{u}_n = 1 - \sum_{1}^{n-1} u_i$. The following proposition describes a family of space-time harmonic functions.

PROPOSITION 3.1. Let $v \in \mathbb{R}^n$ be fixed and let $\varphi(v, \cdot)$ be any primitive of

$$\sum_{i=1}^{n} \left(\mu_i \frac{\phi_i(v,\cdot)}{1 + \phi_i(v,\cdot)} - \lambda_i \phi_i(v,\cdot) \right)$$

on any open subset V of $\{t \ge 0 : 1 + \phi_i(v, t) \ne 0 \text{ for } i = 1, \dots, n\}$. The function

$$h_v(t,x) = e^{\varphi(v,t)} \prod_{i=1}^n (1 + \phi_i(v,t))^{x_i}, \quad t \in V, x \in \mathbb{N}^n,$$

is space-time harmonic with respect to Ω in the domain $V \times \mathbb{N}^{*n}$.

PROOF. It must be shown that $\partial h_v(t,x)/\partial t + \Omega(h_v(t,\cdot))(x) = 0$ on the above domain. For $x \in \mathbb{N}^{*n}$ and $t \in V$, $h_v(t,x) \neq 0$, and one easily computes:

$$\frac{1}{h_v(t,x)}\frac{\partial h_v}{\partial t}(t,x) = \frac{\partial \varphi}{\partial t}(v,t) + \sum_{i=1}^n x_i \frac{\partial \phi_i(v,t)/\partial t}{1 + \phi_i(v,t)}$$

and

$$\frac{1}{h_v(t,x)} \Omega(h_v(t,\cdot))(x) = \sum_{i=1}^n \lambda_i \phi_i(v,t) - \sum_{i=1}^n \mu_i \frac{\phi_i(v,t)}{1 + \phi_i(v,t)} + \sum_{1 \le i \ne j \le n} x_i q_{ij} \frac{\phi_j(v,t) - \phi_i(v,t)}{1 + \phi_i(v,t)}.$$

The last term in the right-hand side is equal to

$$\sum_{i=1}^{n} \frac{x_i}{1 + \phi_i(v, t)} \left(Q\phi(v, t) \right)_i.$$

By definition ϕ satisfies $\partial \phi(v,t)/\partial t = -Q\phi(v,t)$ and the result follows.

REMARK 3.1. The product form of these space-time harmonic functions is quite similar to that of the harmonic functions introduced in Fricker [FRT99] for the $M/M/\infty$ queue.

In addition, it is easily checked that, choosing v = (u - 1, ..., u - 1) for some $u \neq 0$, so that v is some eigenvector of P_t , $t \in \mathbb{R}$, associated to eigenvalue 1, yields $h_v(t, X(t)) = u^{L(t)} e^{[\lambda(1-u)+\mu(1-1/u)]t}$, which is the martingale associated to an M/M/1 queue L with arrival rate λ and service rate μ (see for example Robert [**Rob03**]).

Starting from $h_v(t, x)$, two steps lead to J_{α} : (i) integration of $h_v(t, x)$ over v against some function f(v) on a suitable time-dependent domain $\mathcal{D}(t)$; (ii) change of variables. These two steps are detailed and justified in the Appendix A, yielding the following family of local martingales:

THEOREM 3.1. There exist two positive, continuous, bounded functions F and Gon $\mathring{\mathcal{P}}$ such that for any $\alpha > 0$, $u \mapsto F(\tilde{u})^{\alpha-1}$ is integrable on S and $(J_{\alpha}(t \wedge \mathcal{T}_0), t \geq 0)$ is a nonnegative local martingale, where $J_{\alpha}(t)$ is defined for $\alpha > 0$ and $t \ge 0$ by:

$$J_{\alpha}(t) = e^{-\alpha\theta t} \int_{\mathcal{S}} \prod_{i=1}^{n} \left(\frac{\tilde{u}_{i}}{\pi_{i}}\right)^{X_{i}(t)} G(\tilde{u}) F(\tilde{u})^{\alpha-1} du,$$

or equivalently:

(2)
$$J_{\alpha}(t) = e^{-\alpha\theta t} \int_{\mathcal{S}} e^{L(t)(H(t) - H(\chi(t), \tilde{u}))} G(\tilde{u}) F(\tilde{u})^{\alpha - 1} du.$$

Moreover, F satisfies

(3)
$$\sup_{0<\alpha\leq 1} \left(\alpha^n \int_{\mathcal{S}} F(\tilde{u})^{\alpha-1} du\right) < +\infty.$$

The advantage of $J_{\alpha}(t)$ (as compared to $h_v(t, X(t))$), is that the dependence in time is there splitted into two factors: $e^{-\alpha\theta t}$ is a direct function of time, and the integral is a function of the state of the system at time t, X(t) or equivalently $(L(t), \chi(t))$.

The next proposition gives the fundamental estimate obtained through optional stopping and used several times throughout the chapter.

PROPOSITION 3.2. For any δ such that $0 < \delta < \varepsilon_0$, where ε_0 is given by Lemma 2.2, there exists some constant C_{δ} such that

$$\mathbb{E}_x\left(e^{-\alpha\theta T_H^{\varepsilon}}; L(T_H^{\varepsilon}) \ge \ell\right) \le C_\delta \,\alpha^{-n} e^{|x|H(x/|x|,\pi) - (\varepsilon - \delta)\ell}$$

holds for any initial state $x \in \mathbb{N}^n$ and any $\varepsilon \in]\delta, \varepsilon_0[, \ell > 0 \text{ and } \alpha \in]0, 1].$

Proposition 3.2 is derived from the two following lemmas by choosing $T = T_H^{\varepsilon}$ (so that, by Lemma 2.2, $T \wedge T_0 = T$ when $\varepsilon < \varepsilon_0$). Note that only Lemma 3.1 uses the fact that J_{α} is a local martingale, whereas Lemma 3.2 stems directly from the expression of J_{α} provided by (2).

LEMMA 3.1. There exists some constant $C_3 > 0$ such that, for any $\alpha \in]0,1]$, any initial state $x \in \mathbb{N}^n$ and any stopping time T, the following inequality holds:

$$\mathbb{E}_x\left[J_\alpha(T \wedge \mathcal{T}_0)\right] \le C_3 \alpha^{-n} e^{|x|H(x/|x|,\pi)}.$$

PROOF. Fix $\alpha \in [0, 1]$ and $x \in \mathbb{N}^n$. Since $J_{\alpha}(\cdot \wedge \mathcal{T}_0)$ is a nonnegative local martingale, it is a supermartingale, and so is $(J_{\alpha}(t \wedge T \wedge \mathcal{T}_0), t \geq 0)$ by Doob's optional stopping theorem. In particular, for any $t \geq 0$:

$$\mathbb{E}_x \big[J_\alpha(0) \big] \ge \mathbb{E}_x \big[J_\alpha(t \wedge T \wedge \mathcal{T}_0) \big],$$

and Fatou's lemma gives:

$$\mathbb{E}_x \big[J_\alpha(0) \big] \ge \liminf_{t \to +\infty} \mathbb{E}_x \big[J_\alpha(t \wedge T \wedge \mathcal{T}_0) \big] \ge \mathbb{E}_x \big[\liminf_{t \to +\infty} J_\alpha(t \wedge T \wedge \mathcal{T}_0) \big] = \mathbb{E}_x \big[J_\alpha(T \wedge \mathcal{T}_0) \big],$$

(where $J_{\alpha}(T \wedge \mathcal{T}_0)$ makes sense a.s. when $T \wedge \mathcal{T}_0 = +\infty$ since any nonnegative supermartingale almost surely converges to some variable at infinity).

From the definition of J_{α} given by (2), using $e^{-y} \leq 1$ for $y \geq 0$, one gets

$$\mathbb{E}_{x}\left[J_{\alpha}(0)\right] \leq \sup_{\hat{\mathcal{P}}}(G) e^{|x|H(x/|x|,\pi)} \int_{\mathcal{S}} F(\tilde{u})^{\alpha-1} du \leq C_{3} e^{|x|H(x/|x|,\pi)} \alpha^{-n}$$

where $C_3 = \sup_{\tilde{\mathcal{P}}}(G) \sup_{0 < \alpha \leq 1} \left(\alpha^n \int_{\mathcal{S}} F(\tilde{u})^{\alpha - 1} du \right)$ is finite by (3), which proves the lemma.

LEMMA 3.2. For any positive δ , there exists some positive constant B_{δ} such that the following implication holds for any $\alpha \in [0, 1]$, $\ell > 0$, $\varepsilon > \delta$ and $t \ge 0$:

$$L(t) \ge \ell \text{ and } H(t) \ge \varepsilon \Longrightarrow J_{\alpha}(t) \ge B_{\delta} \cdot e^{-\alpha \theta t + (\varepsilon - \delta)\ell}.$$

PROOF. Fix $\varepsilon > \delta$, $\alpha \in]0, 1]$, $\ell > 0$ and $t \ge 0$. A lower bound on the integral part of (2) is obtained when $L(t) \ge \ell$ and $H(t) \ge \varepsilon$. For $v \in \mathcal{P}$, define the set $\mathcal{S}_{\delta}(v) \subset \mathcal{S}$ by

$$S_{\delta}(v) = \{ u \in S : H(v, \tilde{u}) \le \delta \}.$$

If $H(t) \geq \varepsilon$ and $L(t) \geq \ell$, then

$$\int_{\mathcal{S}} e^{L(t)(H(t) - H(\chi(t), \tilde{u}))} G(\tilde{u}) F(\tilde{u})^{\alpha - 1} du \ge \beta \, e^{\ell(\varepsilon - \delta)} \int_{\mathcal{S}_{\delta}(\chi(t))} G(\tilde{u}) du,$$

where $\beta = \min\{(\sup_{\mathcal{P}} F)^{-1}, 1\}$. Indeed, α being smaller than 1, β is a lower bound for $F(\tilde{u})^{\alpha-1}$ on \mathcal{S} . Consider now the function $\Phi_{\delta} : \mathcal{P} \to \mathbb{R}^+$ defined by

$$\Phi_{\delta}(v) = \int_{\mathcal{S}_{\delta}(v)} G(\tilde{u}) du.$$

Since G is bounded, Φ_{δ} can be shown to be continuous (using for example Lebesgue's theorem). Moreover, $\Phi_{\delta}(v) > 0$ for any $v \in \mathcal{P}$ (because G > 0 and the interior of $\mathcal{S}_{\delta}(v)$ is not empty), and since \mathcal{P} is compact, $\inf_{\mathcal{P}} \Phi_{\delta} > 0$. Setting $B_{\delta} = \beta \inf_{\mathcal{P}} \Phi_{\delta}$ achieves the proof.

4. Two Key Representations

The Markov process $(X(t), t \ge 0)$ with infinitesimal generator Ω defined by (1) can be seen as a particle system involving three types of transitions: births, deaths and migrations of particles from one site to another. The main purpose of this section is to show that X can be decomposed into the difference of two pure birth and migration processes, up to some reflection term (Theorem 4.1). A simpler result (Proposition 4.1) tells that, as long as X does not hit the axis, the process L of the total number of particles just behaves as a random walk (or equivalently as an M/M/1 queue). Finally, a representation of process X involving labelled particles is given in the case of null death rates.

Theorem 4.1, together with the latter representation, will be crucial for describing the unstable regime in Section 6, while Proposition 4.1 will be repeatedly used in the study of both the super and subcritical regimes.

The idea for decomposing X is the following: when $\mu = 0$, the system consists of immortal particles generated at rate λ and performing independent Markov trajectories. Introducing a death procedure, i.e., some positive μ , amounts to eliminating particles (at rate μ_i at site *i*) if possible. Up to some correction due to the fact that no death can actually occur at an empty site, this is equivalent to subtracting some analogous process with birth rates μ_i $(1 \le i \le n)$, zero death rates and migration rate matrix Q.

This can be formalized by introducing an enlarged Markov process involving three types of particles. Define (X, Y, Z) as a Markov process in \mathbb{N}^{3n} with generator Γ characterized by the following transitions and rates: for any $(x, y, z) \in \mathbb{N}^{3n}$ and $i, j \in \{1, \ldots, n\}$ such that $i \neq j$,

$$(x, y, z) \longrightarrow \begin{cases} (x + e_i, y, z) & \text{at rate} \quad \lambda_i \\ (x - e_i, y + e_i, z) & \mu_i \mathbb{1}_{\{x_i \ge 1\}} \\ (x, y, z + e_i) & \mu_i \mathbb{1}_{\{x_i = 0\}} \\ (x - e_i + e_j, y, z) & q_{ij}x_i \\ (x, y - e_i + e_j, z) & q_{ij}y_i \\ (x, y, z - e_i + e_j) & q_{ij}z_i \end{cases}$$

The process X keeps track of the "real" particles, Y of the killed ones and Z of virtual particles generated at some site when no particle has been found to be killed.

It is clear from these transitions and rates that, indexing generator Ω by its birth and death rate vectors: $\underline{\lambda} = (\lambda_i, 1 \leq i \leq n)$ and $\underline{\mu} = (\mu_i, 1 \leq i \leq n)$ $(\underline{\lambda}, \underline{\mu} \in [0, +\infty[^n)$ and denoting by $\underline{0}$ the null vector in \mathbb{R}^n :

(i) X is a Markov process in \mathbb{N}^n with generator $\Omega_{\underline{\lambda},\underline{\mu}}$,

(ii) X + Y is also Markov in \mathbb{N}^n , with generator $\Omega_{\underline{\lambda},\underline{0}}$,

(iii) Y + Z is Markov in \mathbb{N}^n with generator $\Omega_{\underline{\mu},\underline{0}}$,

(iv) |X + Y| - (|X(0) + Y(0)|) is some Poisson process with intensity λ ,

(v) |Y + Z| - (|Y(0) + Z(0)|) is some Poisson process with intensity μ ,

(vi) these two Poisson processes are independent.

Now from (i), any process X with generator Ω can be considered as the first component of some Markov process with generator Γ and initial state $(X(0), \underline{0}, \underline{0})$.

The two next results are easily derived from this construction and from remarks (i) to (vi). In order to state the main theorem, it is convenient to index the process X both by its initial state and by its birth and death parameters, writing $X_{\underline{\lambda},\underline{\mu}}^x$ for the process X with initial state $x \in \mathbb{N}^n$, migration rate matrix Q and birth and death parameters $\underline{\lambda} = (\lambda_i, 1 \leq i \leq n)$ and $\underline{\mu} = (\mu_i, 1 \leq i \leq n)$ respectively.

THEOREM 4.1. For any $x \in \mathbb{N}^n$ and $\underline{\lambda}, \underline{\mu} \in [0, +\infty[^n, \text{ there exist versions of } X^x_{\underline{\lambda},\underline{\mu}}, X^x_{\underline{\lambda},\underline{0}} \text{ and } X^{\underline{0}}_{\underline{\mu},\underline{0}} \text{ such that}$

$$X^x_{\underline{\lambda},\underline{\mu}} = X^x_{\underline{\lambda},\underline{0}} - X^{\underline{0}}_{\underline{\mu},\underline{0}} + Z,$$

where Z is an \mathbb{N}^n -valued process such that |Z| is nondecreasing, initially zero, and increases only at times when some $X_i(t)$ is zero.

PROOF. Write X = X + Y - (Y + Z) + Z, where $(X(0), Y(0), Z(0)) = (x, \underline{0}, \underline{0})$ and (X, Y, Z) is Markov with generator Γ , so that X is some version of $X^x_{\underline{\lambda},\underline{\mu}}$, and by (ii) and (iii), X + Y is some version of $X^x_{\underline{\lambda},\underline{0}}$ and Y + Z some version of $X^{\underline{0}}_{\underline{\mu},\underline{0}}$. The theorem is proved, since |Z| has the stated properties as can be seen on Γ .

REMARK 4.1. The process Z in Theorem 4.1 appears as a reflection term: it guarantees that X stays nonnegative, compensating by adding some virtual particle for a jump of $X_{\mu,0}^{\underline{0}}$ that would get some X_i to the value -1.

However, contrary to usual multidimensional Skorokhod reflection terms, here, due to the movements of particles, components Z_i 's are not necessarily nondecreasing in time: only their sum is; also Z_i can increase at times when X_i is not zero.

Theorem 4.1 and its proof, together with properties (iv), (v) and (vi), give the following proposition, which constitutes one of the key ingredients for deriving the fluid limits in Sections 6 and 7.

PROPOSITION 4.1. For all $t \leq T_0$, the following equality holds:

$$L(t) = L(0) + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t),$$

where \mathcal{N}_{λ} and \mathcal{N}_{μ} are independent Poisson processes with respective intensities λ and μ . Moreover, $L(t) \geq L(0) + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t)$ holds for any $t \geq 0$.

We conclude this section with a representation of process $X_{\lambda,0}^x$ that will notably be used in Section 6, in conjunction with Theorem 4.1, for analyzing the unstable regime. The process $X_{\underline{\lambda},\underline{0}}^x$ is here obtained as function of a Poisson process with intensity λ and a sequence of Markov processes with infinitesimal generator Q (representing the trajectories of the successively generated particles).

More precisely, $X_{\lambda,0}^x$ admits the following representation:

(4)
$$X_{\underline{\lambda},\underline{0}}^{x}(t) = \Big(\sum_{k\geq 1} \mathbb{1}_{\{\xi_{k}(t-\sigma_{k})=i, \sigma_{k}\leq t\}}, 1\leq i\leq n\Big), \quad t\geq 0,$$

where

- $\sigma_k = 0$ for $1 \le k \le |x|$,
- $\mathcal{N}_{\lambda} = (\sigma_k, k \ge |x| + 1)$ is a Poisson process with parameter λ ,
- $\xi_k, k \ge 1$, are Markov jump processes in $\{1, \ldots, n\}$ with generator Q and initial distribution

 - $\sum_{i=1}^{n} (\lambda_i / \lambda) \,\delta_i \quad \text{for } k \ge |x| + 1, \\ \delta_i \text{ for } x_i \text{ arbitrarily chosen indices } k \in \{1, \dots, |x|\} \ (1 \le i \le n),$
- \mathcal{N}_{λ} and the $\xi_k, k \geq 1$, are mutually independent.

The vector $(\xi_k, 1 \leq k \leq |x|)$ holds for the trajectories of the initial particles and $(\xi_k, k \ge |x|+1)$ for those of the successive newborn particles; the Poisson process \mathcal{N}_{λ} holds for the global birth process. For $k \ge 1$, particle k is in the system from time σ_k , with $\sigma_k = 0$ for $k \leq |x|$.

Similarly as in the previous construction, a formal proof of Equation (4) can be provided by constructing $X_{\underline{\lambda},\underline{0}}^x$ as function of a more complete process (that also contains \mathcal{N}_{λ} and the ξ_k 's, $k \geq 1$, characterized through its infinitesimal generator and describing the list of current positions of particles present in the system, ordered according to their birth rank.

5. The Space Renormalized Process

The stability property of the system for $\lambda < \mu$ will be derived in Section 7 from a fluid scaling analysis, that is, from the study of the space-time renormalized process

$$\overline{X}^{x}(t) = \frac{X^{x}(|x|t)}{|x|}, \quad t \ge 0,$$

as |x| goes to infinity, where X^x is the Markov process X initiated at x. It will be underlain by the M/M/1 behavior of the total occupancy process L (only valid as long as no X_i is zero, hence the intricacies of the analysis).

The particular behavior of \overline{X}^x at $t = 0^+$ will result from the short term behavior of the only space renormalized process \hat{X} , defined as the the family of processes

$$\widehat{X}^x(t) = \frac{X^x(t)}{|x|}, \quad t \ge 0, \quad \text{ for } x \in \mathbb{N}^n \setminus \{0\}.$$

The simpler notation $\widehat{X}(t)$, where $\widehat{X}(t) = X(t)/|X(0)|$, will also be used in situations where |X(0)| is clearly non zero.

As highlighted in the introduction, this scaling is natural and analogous to the Kelly scaling for the $M/M/\infty$ queue. This analogy appears in Proposition 5.3 below, that states convergence of \widehat{X}^x as $|x| \to +\infty$ to some dynamical system having π as its limiting point. In particular, for large |x|, \widehat{X}^x reaches any neighborhood of π in a quasi-deterministic finite time. And this will show (Sections 6 and 7) that asymptotically, \overline{X}^x is instantaneously at π .

The results of this section are quite standard, essentially based on law of large numbers principles. The simple underlying idea is that, as far as X^x is only observed over a finite time window, since the number |x| of initial particles goes to infinity while the numbers of births and deaths within the given window remain of the order of 1 (time is not rescaled here), the initial particles asymptotically dominate the system and mostly stay alive all along the time window, thus behaving as |x|independent Markov processes with generator Q.

For the same reasons, the process \widehat{X} is not different, in the limit $|x| \to +\infty$, from the process $\chi = X/L$ of the spatial distribution of particles: The same convergence results hold for both processes; once proved for \widehat{X} , they easily extend to χ .

Formalizing the above argument, the following coupling is intuitively clear. It compares the general model to the "closed" one (with no births nor deaths, but only initial particles). As in Section 4, generator Ω is indexed by its birth and death parameters $\underline{\lambda}$ and μ .

LEMMA 5.1. For any $x \in \mathbb{N}^n$, there exists a coupling between the process X^x with initial state x and generator $\Omega_{\underline{\lambda},\mu}$, and the process U^x with initial state x and generator $\Omega_{\underline{0},\underline{0}}$, such that, for $t \geq 0$ and $i = 1, \ldots, n$:

$$U_i^x(t) - \mathcal{N}_{\mu}(t) \le X_i^x(t) \le U_i^x(t) + \mathcal{N}_{\lambda}(t),$$

where \mathcal{N}_{λ} and \mathcal{N}_{μ} are two Poisson processes with respective parameters λ and μ . The process X^x moreover satisfies

$$|x| - \mathcal{N}_{\mu}(t) \le |X^{x}(t)| \le |x| + \mathcal{N}_{\lambda}(t), \quad t \ge 0.$$

PROOF. The case $\mu = 0$ is a straightforward consequence of the representation (4) of X from Section 4. Indeed if $\mu = 0$, then (4) gives for $1 \le i \le n$,

$$U_i^x(t) \le X_i^x(t) = \sum_{1 \le k \le |x|} \mathbb{1}_{\{\xi_k(t)=i\}} + \sum_{k \ge |x|+1} \mathbb{1}_{\{\xi_k(t-\sigma_k)=i, \sigma_k \le t\}} \le U_i^x(t) + \mathcal{N}_\lambda(t),$$

where $U^x(t)$ is constructed as $\sum_{k=1}^{|x|} \mathbb{1}_{\{\xi_k(t)=i\}}$ and $\mathbb{N}_{\lambda} = (\sigma_k, k \ge |x|+1)$. One moreover gets $|U^x(t)| \le |X^x(t)| \le |U^x(t)| + \mathcal{N}_{\lambda}(t)$ by summing up over *i* the previous first inequalities. The lemma is proved in this case since $|U^x(t)| = |x|$ for any $t \ge 0$.

The general case is then derived, using the first part of Section 4. Indeed, consider X^x as the first component of a random process (X^x, Y, Z) in \mathbb{N}^{3n} such that Y(0) = Z(0) = 0, $X^x + Y$ is some process with generator $\Omega_{\lambda,0}$ and |Y + Z| is some Poisson process \mathcal{N}_{μ} with intensity μ . The first part of the proof then applies to $X^x + Y$ and gives, for $t \ge 0$:

$$U^{x}(t) \leq X^{x}(t) + Y(t) \leq U^{x}(t) + \mathcal{N}_{\lambda}(t)$$

componentwise, as well as $|x| \leq |X^x(t) + Y(t)| \leq |x| + \mathcal{N}_{\lambda}(t)$.

The lemma follows by noticing that

$$0 \le Y_i(t) \le |Y(t)| \le |Y(t) + Z(t)| = \mathcal{N}_{\mu}(t)$$

holds for any $1 \leq i \leq n$.

The two main results of this section concern the hitting time of some neighborhood of π by the space renormalized process \widehat{X} : namely, for any positive δ ,

$$\widehat{T}_{\delta} = \inf \left\{ t \ge 0 : \left\| \widehat{X}(t) - \pi \right\| \le \delta \right\}.$$

Recall that the analogous time with χ in place of \widehat{X} is denoted by T_{δ} .

PROPOSITION 5.1. For any positive δ , there exists some deterministic time $t_{\delta} \geq 0$ such that

$$\lim_{|x|\to+\infty} \mathbb{P}_x(\widehat{T}_\delta > t_\delta) = 0.$$

The same result holds for the stopping time T_{δ} .

PROOF. We refer to the proof of Proposition 5.2 below. Proposition 5.1 is obtained in the same way, just changing δ_N , s_N and t_N into δ , $s = -(1/\eta) \log(\delta/2B)$ and $t = -1/\eta \log(\delta/4B)$.

The following more accurate result will be required for analyzing the subcritical case $\lambda < \mu$.

PROPOSITION 5.2. There exist two positive constants A and η such that, for any sequence of positive numbers $(\delta_N, N \ge 1)$ satisfying

$$\lim_{N \to +\infty} \delta_N = 0 \text{ and } \lim_{N \to +\infty} \delta_N \sqrt{N} = +\infty,$$

then

$$\lim_{N \to +\infty} \left[\max_{x \in \mathbb{N}^n : |x| = N} \mathbb{P}_x \left(\widehat{T}_{\delta_N} > t_N \right) \right] = 0, \quad \text{where } t_N = -\frac{1}{\eta} \log \frac{\delta_N}{A}.$$

The same result holds for the stopping time T_{δ_N} .

PROOF. First consider a closed system, i.e., assume $\lambda = \mu = 0$; the general case will then be deduced from Lemma 5.1. As in Lemma 5.1, let U^x be the closed process with initial state $x \in \mathbb{N}^n$, where |x| = N. In this case (4) becomes

$$U_i^x(t) = \sum_{k=1}^N \mathbb{1}_{\{\xi_k(t)=i\}}, \quad 1 \le i \le n, \ t \ge 0,$$

where ξ_k , $1 \le k \le N$, are independent Markov processes with the same generator Q and different initial conditions: for any $i, 1 \le i \le n, \xi_k(0) = i$ for x_i of the N indices k = 1, ..., N.

As introduced in Section 3, let $(P_t, t \ge 0)$ denote the transition semi-group associated to Q. The exponentially fast convergence of any irreducible finite state

space Markov semi-group to its stationary distribution, tells existence of B>0 and $\eta>0$ such that

$$\max_{1 \le i,j \le n} |P_t(j,i) - \pi_i| \le Be^{-\eta t}, \quad t \ge 0.$$

In particular, for $s_N = -(1/\eta) \log(\delta_N/2B)$,

$$\max_{1 \le i,j \le n} |P_{s_N}(j,i) - \pi_i| \le \delta_N/2$$

The outline of the proof for the closed case is the following: At time s_N , all trajectories ξ_k , $1 \leq k \leq N$, are very close to π in distribution (by the order of δ_N). Since $\hat{U}^x(t)$ represents the empirical distribution of the N particles at time t, the law of large numbers shows that for large N, $\hat{U}^x(s_N)$ is also close to π (by the same order), because δ_N tends to 0 not too fast

Precisely, for any $N \ge 1$ and $x \in \mathbb{N}^n$ such that |x| = N:

$$\left\|\mathbb{E}\left(\widehat{U}^{x}(s_{N})\right) - \pi\right\| = \left\|\mathbb{E}\left(\frac{U^{x}(s_{N})}{N}\right) - \pi\right\| = \left\|\frac{1}{N}\sum_{k=1}^{N}\left(\mathbb{P}(\xi_{k}(s_{N}) = \cdot) - \pi\right)\right\| \le \frac{\delta_{N}}{2}.$$

Thus, for any $N \ge 1$, using Chebychev's inequality for the last step:

$$\mathbb{P}\left(\|\widehat{U}^{x}(s_{N})-\pi\|>\delta_{N}\right) \leq \mathbb{P}\left(\left\|\widehat{U}^{x}(s_{N})-\mathbb{E}\left(\widehat{U}^{x}(s_{N})\right)\right\|>\frac{\delta_{N}}{2}\right)$$
$$\leq \sum_{i=1}^{n} \mathbb{P}\left(\left|\left.\widehat{U}^{x}_{i}(s_{N})-\mathbb{E}\left(\widehat{U}^{x}_{i}(s_{N})\right)\right|>\frac{\delta_{N}}{2}\right) \leq \sum_{i=1}^{n} \frac{\mathbb{V}\mathrm{ar}\left(U^{x}_{i}(s_{N})\right)}{\delta_{N}^{2}N^{2}/4}.$$

Independence of the processes $(\xi_k, 1 \le k \le N)$ yields

$$\mathbb{V}\mathrm{ar}\left(U_{i}^{x}(s_{N})\right) = \sum_{k=1}^{N} \mathbb{V}\mathrm{ar}\left(\mathbb{1}_{\left\{\xi_{k}(s_{N})=i\right\}}\right) \leq \frac{N}{4}$$

(bounding the variance of any Bernoulli random variable by 1/4). Finally

(5)
$$\max_{x \in \mathbb{N}^n : |x|=N} \mathbb{P}\left(\|\widehat{U}^x(s_N) - \pi\| > \delta_N\right) \le \frac{n}{\delta_N^2 N}.$$

Now consider the process X^x associated to any family $(\lambda_i, \mu_i, 1 \leq i \leq n)$ of parameters and any initial state x such that |x| = N. Still denote by U^x the associated closed process with the same initial state x.

Define $t_N = -(1/\eta) \log(\delta_N/4B)$. The first part of Lemma 5.1 implies that, for any $N \ge 1$,

$$\left\|\widehat{X}^{x}(t_{N})-\pi\right\| \leq \left\|\widehat{U}^{x}(t_{N})-\pi\right\| + \left\|\frac{1}{N}\left(\mathcal{N}_{\lambda}(t_{N})+\mathcal{N}_{\mu}(t_{N})\right)\right\|$$

so that

$$\mathbb{P}_{x}(\widehat{T}_{\delta_{N}} > t_{N}) \leq \mathbb{P}\left(\|\widehat{U}^{x}(t_{N}) - \pi\| > \frac{\delta_{N}}{2}\right) + \mathbb{P}\left(\|\mathcal{N}_{\lambda}(t_{N}) + \mathcal{N}_{\mu}(t_{N})\| > \frac{N\delta_{N}}{2}\right).$$

By (5) the first term tends to zero uniformly in x as N goes to infinity, since t_N is associated to $\delta_N/2$ in the same way as s_N was to δ_N . The second one is also easily shown to converge to zero, using Chebychev's inequality for the Poisson variable $\mathcal{N}_{\lambda}(t_N) + \mathcal{N}_{\mu}(t_N)$, together with the relation $\delta_N \sqrt{N} \gg 1$ that implies $N\delta_N \gg \sqrt{N} \gg 1/\delta_N \gg t_N$.

The first part of the proposition is thus proved with A = 4B.

Using the last assertion of Lemma 5.1, it is not difficult to show that the same result holds for T_{δ_N} .

We finally just mention for the sake of completeness (it will not be used in the sequel) the following result that describes the asymptotic dynamics, as $|x| \to +\infty$, of the empirical distribution of the particles: it evolves as the distribution, as function of time, of a Markov process with generator Q.

Not surprisingly, this can be proved using the same standard arguments as for studying the $M/M/\infty$ queue under the Kelly scaling (see Robert [**Rob03**]).

PROPOSITION 5.3. Consider the processes $(\widehat{X}^{x_N}(t), t \ge 0)$ associated to some sequence $(x_N, N \ge 1)$ of initial states satisfying: $\lim_{N \to +\infty} \frac{x_N}{N} = \rho$ for some $\rho \in \mathcal{P}$.

For any T > 0, as $N \to +\infty$, $(\widehat{X}^{x_N}(t), t \ge 0)$ converges in distribution with respect to the uniform norm topology on [0, T], to the deterministic trajectory:

$$\rho(t) = \rho P_t.$$

In other words, for any positive δ :

$$\lim_{N \to +\infty} \mathbb{P}\Big(\sup_{0 \le t \le T} \|\widehat{X}^{x_N}(t) - \rho P_t\| > \delta\Big) = 0.$$

The same convergence holds for the corresponding processes $(\chi^{x_N}(t), t \ge 0), N \ge 1$.

6. The Supercritical Regime

This section deals with the supercritical regime $\lambda > \mu$. As the next proposition shows, the unstability of the system is straightforward in this case. Theorem 6.1 establishes an almost sure result describing the long term behavior, and Theorem 6.2 presents a surprising phenomenon.

PROPOSITION 6.1. When $\lambda > \mu$ the process X is not ergodic.

PROOF. Just remark, using Proposition 4.1, that if $x \in \mathbb{N}^n$ is the initial state:

$$L(t) \ge |x| + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t).$$

Hence for any initial state, L(t) almost surely goes to $+\infty$ as t tends to $+\infty$.

The following theorem gives an almost sure description of the divergence of X(t) for t large. Among other arguments, the proof makes use for the first time of the martingale estimate provided by Proposition 3.2, and involves the representations of X given in Section 4.

THEOREM 6.1. Assume $\lambda > \mu$. Then, for any initial state $x \in \mathbb{N}^n$, the following convergence holds almost surely:

$$\lim_{t \to +\infty} \frac{X^x(t)}{t} = (\lambda - \mu)\pi.$$

REMARK 6.1. This theorem has a double meaning: it tells almost sure convergence both of $\chi(t)$ to π and of L(t)/t to $\lambda - \mu$ as $t \to +\infty$. PROOF. Assume the theorem is true when $\mu = 0$. Then, using the notations of Theorem 4.1, $t^{-1}(X_{\underline{\lambda},\underline{0}}^x(t) - X_{\underline{\mu},\underline{0}}^{\underline{0}}(t))$ converges a.s. to $(\lambda - \mu)\pi$ and the componentwise inequality $X_{\underline{\lambda},\underline{\mu}}^x \geq X_{\underline{\lambda},\underline{0}}^x - X_{\underline{\mu},\underline{0}}^{\underline{0}}$ derived from Theorem 4.1, implies that each $X_i^x(t)$ tends to infinity almost surely as t goes to infinity.

As a consequence, since |Z(t)| can increase only when some $X_i(t)$ is zero, then, with probability 1, $\lim_{t\to+\infty} |Z(t)|$ is finite and $\lim_{t\to+\infty} Z(t)/t = 0$, so that:

$$\lim_{t \to +\infty} \frac{X^x(t)}{t} = \lim_{t \to +\infty} \frac{X^x_{\underline{\lambda},\underline{0}}(t) - X^{\underline{0}}_{\underline{\mu},\underline{0}}(t)}{t} = (\lambda - \mu)\tau$$

holds almost surely, which is the stated result.

The theorem must now be proved in the case where $\mu = 0$. In this case with no deaths, using representation (4), the process X^x splits into two (independent) processes: $X^x = U^x + X^0$, where U^x is associated to a "closed" system with |x|particles moving independently, and X^0 has no initial particles, birth rates $\underline{\lambda}$ and null death rates. Then $t^{-1}U^x(t)$ obviously tends to zero almost surely as t tends to infinity, and all is left to show is that $t^{-1}X^0(t)$ converges almost surely to $\lambda \pi$.

So, dropping for simplicity the superscript $\underline{0}$, consider the process X with initial state $\underline{0}$, birth rates $\underline{\lambda}$ and null death rates. Equation (4) here becomes, for $t \ge 0$:

$$X_{i}(t) = \sum_{k=1}^{\mathcal{N}_{\lambda}(t)} \mathbb{1}_{\{\xi_{k}(t-\sigma_{k})=i\}} , \quad 1 \le i \le n,$$

where $(\xi_k, k \ge 1)$, have initial distribution $\sum_{i=1}^n (\lambda_i / \lambda) \delta_i$.

It will first be shown that the analysis can be reduced to the case of stationary trajectories (i.e., the case when $\lambda_i/\lambda = \pi_i$ for $1 \leq i \leq n$) by using a coupling argument.

Indeed, associate to each ξ_k a stationary process ξ'_k with the same generator, such that $((\xi_k, \xi'_k), k \ge 1)$ is a sequence of independent processes in $\{1, \ldots, n\}^2$, and, for $k \ge 1$, ξ_k , ξ'_k are coupled in the classical following way: ξ_k and ξ'_k are independent until the first time T_k when they meet, and after that stay equal for ever. Recall that the "coupling times" T_k , $k \ge 1$, are integrable. Moreover assume the $(\xi_k, \xi'_k), k \ge 1$, independent from \mathcal{N}_{λ} .

Define the process $(X'(t), t \ge 0)$ on \mathbb{N}^n analogously to X, with the same \mathcal{N}_{λ} , but with ξ'_k in place of ξ_k $(k \ge 1)$. Then, for each $i \in \{1, \ldots, n\}$

$$|X_i(t) - X_i'(t)| = \left|\sum_{k=1}^{\mathcal{N}_{\lambda}(t)} \left(\mathbbm{1}_{\{\xi_k(t-\sigma_k)=i\}} - \mathbbm{1}_{\{\xi_k'(t-\sigma_k)=i\}}\right)\right| \le \sum_{k=1}^{\mathcal{N}_{\lambda}(t)} \mathbbm{1}_{\{T_k > t-\sigma_k\}}.$$

Define the right hand side of the above equation as A(t): then for any $t \ge 0$, A(t) is exactly the number of customers at time t in an $M/G/\infty$ queue with no customer at time 0, arrival process \mathcal{N}_{λ} , and services given by the i.i.d. integrable variables $T_k, k \ge 1$. It is easily proved that A(t)/t converges almost surely to zero as t tends to infinity. It is then enough to prove a.s. convergence of process X' to $\lambda \pi$, and so we assume from now on that $(\xi_k, k \ge 1)$ are stationary.

Since $L(t)/t = \mathcal{N}_{\lambda}(t)/t$ converges a.s. to λ as t tends to infinity, the problem is equivalent to proving that $\chi(t)$ converges almost surely to π , i.e., by Lemma 2.1 that:

$$\forall \varepsilon > 0, \ \mathbb{P}(\exists T < +\infty : \forall t \geq T, H(t) \leq \varepsilon) = 1$$

This will be done using Borel-Cantelli lemma and showing that:

$$\forall \varepsilon > 0, \ \sum_{k=1}^{+\infty} \mathbb{P}(\exists t \in [\sigma_k, \sigma_{k+1}[: H(t) > \varepsilon) < +\infty,$$

Writing, for any fixed ε :

(6)
$$\mathbb{P}\left(\exists t \in [\sigma_k, \sigma_{k+1}[: H(t) > \varepsilon)\right)$$

 $\leq \mathbb{P}\left(H(\sigma_k) > \frac{\varepsilon}{2}\right) + \mathbb{P}\left(H(\sigma_k) \leq \frac{\varepsilon}{2} \text{ and } \exists t \in]\sigma_k, \sigma_{k+1}[: H(t) > \varepsilon\right),$

we will show that both series associated to both terms in the right hand side converge for ε sufficiently small (which is enough by monotonicity of the left hand side of (6)).

Let us begin with the first term. Note that for $k \ge 1$, $\chi(\sigma_k) = X(\sigma_k)/k$. Then due to Lemma 2.1, it is enough to show that, for small ε and any $i \in \{1, \ldots, n\}$,

(7)
$$\sum_{k=1}^{+\infty} \mathbb{P}\left(\left|\frac{X_i(\sigma_k)}{k} - \pi_i\right| > \varepsilon\right) < +\infty.$$

This is obtained by using Chernoff's inequality, that we recall in Lemma 6.1.

LEMMA 6.1. (Chernoff's inequality) Let $Z_h, 1 \le h \le k$, be k independent random variables such that $|Z_h| \le 1$ and $\mathbb{E}(Z_h) = 0$ for $1 \le h \le k$. The following bound holds for any $\eta \in [0, 2\sigma]$, where $\sigma^2 = \mathbb{V}\mathrm{ar}\left(\sum_{h=1}^k Z_h\right)$:

$$\mathbb{P}\left(\left|\sum_{h=1}^{k} Z_{h}\right| \ge \eta\sigma\right) \le 2e^{-\eta^{2}/4}$$

Write for $1 \leq i \leq n$

$$\frac{X_i(\sigma_k)}{k} - \pi_i = \frac{1}{k} \sum_{h=1}^k Z_{k,h}^{(i)} \text{ with } Z_{k,h}^{(i)} = \mathbb{1}_{\{\xi_h(\sigma_k - \sigma_h) = i\}} - \pi_i.$$

Since $\xi_h, h \ge 1$, are stationary, then for each fixed $i \in \{1, \ldots, n\}$ and $k \ge 1$, the k variables $Z_{k,h}^{(i)}, 1 \le h \le k$, are i.i.d. centered random variables, bounded by 1 in modulus. (Notice that independence is only true in this stationary case).

We can thus apply Chernoff's inequality, which gives, for each fixed k and i:

$$\mathbb{P}\left(\left|\frac{X_i(\sigma_k)}{k} - \pi_i\right| > \varepsilon\right) = \mathbb{P}\left(\left|\sum_{h=1}^k Z_{k,h}^{(i)}\right| > k\varepsilon\right) \le 2e^{-\frac{\varepsilon^2 k}{4v_i}},$$

if $\varepsilon \leq 2v_i$, where $v_i = \pi_i(1 - \pi_i)$ is the common variance of the variables $Z_{k,h}^{(i)}$. Property (7) is then proved (for small ε , hence for any ε by monotonicity).

Now it must be shown that the second term in the right hand side of (6) is summable as well for ε small enough. Here, the stationarity of the movements will play no special role.

By definition of σ_k , $\chi(t) = X(t)/k$ for any $t \in [\sigma_k, \sigma_{k+1}]$. Moreover, σ_k is a stopping time for the Markov process $(X(t), t \ge 0)$, because it is the first time
when L(t) = k. Hence the strong Markov property yields

$$\mathbb{P}_{\underline{0}}\left(H(\sigma_k) \leq \frac{\varepsilon}{2} \text{ and } \exists t \in]\sigma_k, \sigma_{k+1}[:H(t) > \varepsilon\right) \leq \max_{\substack{x \in \mathbb{N}^n: |x| = k \text{ and} \\ H(x/|x|, \pi) \leq \varepsilon/2}} \mathbb{P}_x\left(T_H^{\varepsilon} < \sigma_1\right).$$

Clearly, the last event only depends on σ_1 and on the movements of the |x| initial particles, so that by independence of these variable and processes, one obtains, for any $x \in \mathbb{N}^n$:

$$\mathbb{P}_x\left(T_H^{\varepsilon} < \sigma_1\right) = \mathbb{E}_x\left(e^{-\lambda \widetilde{T}_H^{\varepsilon}}\right) \le \mathbb{E}_x\left(e^{-(\lambda \wedge \theta)\widetilde{T}_H^{\varepsilon}}\right),$$

where T_H^{ε} is the first time the entropy associated to the initial particles is larger than ε . Then, using Proposition 3.2 in the case of a closed system with $\delta = \varepsilon/4$, $\alpha = (\lambda/\theta) \wedge 1$ and $\ell = k$ gives:

$$\mathbb{P}_x\left(T_H^{\varepsilon} < \sigma_1\right) \le C_{\varepsilon/4}[(\lambda/\theta) \wedge 1)]^{-n} e^{-\varepsilon k/4},$$

for any $x \in \mathbb{N}^n$ such that |x| = k and $H(x/|x|, \pi) \leq \varepsilon/2$. The second term in (6) is thus summable over k for ε small enough.

Along the preceding proof, we used σ_1 , in the particular case $\mu = 0$, as an asymptotic lower bound (as the initial state grows to infinity) for the exit time of $\chi(t)$ from some neighborhood of π . This is a very crude underestimation, as the following result shows that this exit time is actually infinite with high probability.

THEOREM 6.2. Assume $\lambda > \mu$, and fix δ and ε such that $0 < \delta < \varepsilon < \varepsilon_0$, where ε_0 is given by Lemma 2.2. Consider a sequence $(x_N, N \ge 1)$ with $\lim_{N\to+\infty} |x_N|/N = 1$ and $H(x_N/|x_N|, \pi) \le \delta$. Then:

$$\lim_{N \to +\infty} \mathbb{P}_{x_N} \left(T_H^{\varepsilon} = +\infty \right) = 1.$$

PROOF. By definition of T_H^{ε} , \mathbb{P}_{x_N} ($T_H^{\varepsilon} < +\infty$) = \mathbb{P}_{x_N} ($\exists t \ge 0$: $H(t) \ge \varepsilon$), and so we need to study the behavior of H(t) for all time $t \ge 0$. The idea of the proof is twofold: First, the estimate given by Proposition 3.2 is precise enough to show that T_H^{ε} is much larger than N, say $T_H^{\varepsilon} \ge N^2$. After this time, the initial particles are negligible, and Theorem 6.1 then gives a control on the rest of the trajectory by reducing the problem to the case where the system starts empty. So we use the following decomposition:

$$\mathbb{P}_{x_N}\left(T_H^{\varepsilon} < +\infty\right) \le \mathbb{P}_{x_N}\left(T_H^{\varepsilon} \le N^2\right) + \mathbb{P}_{x_N}\left(\exists t \ge N^2 : H(t) \ge \varepsilon\right).$$

For the first term, Markov's inequality gives

(8)
$$\mathbb{P}_{x_N}\left(T_H^{\varepsilon} \le N^2\right) \le e\mathbb{E}_{x_N}\left(e^{-T_H^{\varepsilon}/N^2}\right)$$

Let $\delta' < \varepsilon - \delta$: by choice of ε and δ , and since $H(x_N/|x_N|, \pi) \leq \delta$, Proposition 3.2 shows that there exists a constant $C_{\delta'}$ such that by choosing $\alpha = 1/(\theta N^2)$, for any N large enough and any ℓ_N ,

$$\mathbb{E}_{x_N}\left(e^{-T_H^{\varepsilon}/N^2}; L(T_H^{\varepsilon}) \ge \ell_N\right) \le C_{\delta'} e^{\delta |x_N| + 2n \log N - (\varepsilon - \delta')\ell_N}.$$

The choice of ℓ_N requires some care: as N grows, it must be both of order $|x_N|$ and smaller than $L(T_H^{\varepsilon})$ with high probability. Since $|x_N| \sim N$, write $|x_N| = N + u_N$ with $u_N = o(N)$, and choose $\ell_N = N - \sqrt{Nv_N}$ with $v_N = |u_N| \lor 1$. With this choice, $\ell_N \sim N$ and $\ell_N - |x_N| \to -\infty$. The first relation implies, since $\varepsilon - \delta' - \delta > 0$,

$$\lim_{N \to +\infty} e^{\delta |x_N| + 2n \log N - (\varepsilon - \delta')\ell_N} = 0.$$

Moreover, since $T_H^{\varepsilon} \leq \mathcal{T}_0$ because $\varepsilon < \varepsilon_0$, Proposition 4.1 implies that $L(T_H^{\varepsilon}) = L(0) + \mathcal{N}_{\lambda}(T_H^{\varepsilon}) - \mathcal{N}_{\mu}(T_H^{\varepsilon})$, hence

$$\mathbb{P}_{x_N} \left(L(T_H^{\varepsilon}) \le \ell_N \right) = \mathbb{P}_{x_N} \left(|x_N| + \mathcal{N}_{\lambda}(T_H^{\varepsilon}) - \mathcal{N}_{\mu}(T_H^{\varepsilon}) \le \ell_N \right)$$
$$\le \mathbb{P} \left(\inf_{t \ge 0} \left(\mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t) \right) \le \ell_N - |x_N| \right)$$

where the last bound vanishes because $\lambda > \mu$, and so $\inf_{t\geq 0} (\mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t))$ is finite with probability one, whereas $\ell_N - |x_N|$ goes to $-\infty$. It results that $\mathbb{P}_{x_N} (T_H^{\varepsilon} \leq N^2)$ goes to 0 thanks to (8) and to the following inequality:

$$\mathbb{E}_{x_N}\left(e^{-T_H^{\varepsilon}/N^2}\right) \le \mathbb{E}_{x_N}\left(e^{-T_H^{\varepsilon}/N^2}; L(T_H^{\varepsilon}) \ge \ell_N\right) + \mathbb{P}_{x_N}\left(L(T_H^{\varepsilon}) \le \ell_N\right),$$

and it has been shown that each term goes to 0.

All is left to prove now is that $\lim_{N\to+\infty} \mathbb{P}_{x_N} \left(\exists t \geq N^2 : H(t) \geq \varepsilon \right) = 0$, or, by Lemma 2.1, that $\mathbb{P}_{x_N} \left(\exists t \geq N^2 : \|\chi(t) - \pi\| \geq \varepsilon \right)$ vanishes. After time N^2 , the initial particles are negligible since a number of new particles of the order of N^2 have arrived. So the behavior of the system will be similar to that of a system starting empty, to which we can apply Theorem 6.1 (since in this case the initial state is fixed).

To formalize this argument, a coupling between the processes X^x and X^0 , for any $x \in \mathbb{N}^n$, is required:

LEMMA 6.2. For any $x, y \in \mathbb{N}^n$ with $x \geq y$ componentwise, it is possible to couple the two processes X^x and X^y in such a way that for any $t \geq 0$, $L^x(t) - L^y(t) \leq |x| - |y|$ and the inequality $X^x(t) \geq X^y(t)$ holds componentwise.

The proof of this lemma is postponed at the end of the current proof. Let X^0 be the process starting empty coupled with X^{x_N} , and let $L^0 = |X^0|$, $L^{x_N} = |X^{x_N}|$, $\chi^0 = X^0/L^0$ and $\chi^{x_N} = X^{x_N}/L^{x_N}$ be the corresponding quantities. The triangular inequality gives

(9)
$$\mathbb{P}(\exists t \ge N^2 : \|\chi^{x_N}(t) - \pi\| \ge \varepsilon) \le \mathbb{P}(\exists t \ge N^2 : \|\chi^{x_N}(t) - \chi^0(t)\| \ge \varepsilon/2)$$
$$+ \mathbb{P}(\exists t \ge N^2 : \|\chi^0(t) - \pi\| \ge \varepsilon/2).$$

Theorem 6.1 states that $\chi^0(t)$ converges to π almost surely, which shows that the last term goes to 0. For the first term, write for each i = 1, ..., n,

$$\chi_i^{x_N}(t) - \chi_i^0(t) = \frac{X_i^{x_N}(t)}{L^{x_N}(t)} - \frac{X_i^0(t)}{L^0(t)} = \frac{\left(X_i^{x_N}(t) - X_i^0(t)\right)L^0(t) - X_i^0(t)\Delta^{x_N}(t)}{L^0(t)\left(\Delta^{x_N}(t) + L^0(t)\right)},$$

where $\Delta^{x_N}(t) = L^{x_N}(t) - L^0(t)$. Lemma 6.2 implies that $|X_i^{x_N}(t) - X_i^0(t)| \leq \Delta^{x_N}(t) \leq |x_N|$, hence, since the function $z \mapsto z/(z+a)$ is decreasing for any $a \geq 0$,

$$|\chi_i^{x_N}(t) - \chi_i^0(t)| \le \frac{2\Delta^{x_N}(t)}{\Delta^{x_N}(t) + L^0(t)} \le \frac{2|x_N|}{|x_N| + L^0(t)} = \frac{2}{1 + L^0(t)/|x_N|}.$$

This yields in turn, using $t \ge N^2$ for the second inequality,

$$\mathbb{P}(\exists t \ge N^2 : \|\chi^{x_N}(t) - \chi^0(t)\| \ge \varepsilon/2) \le \mathbb{P}\left(\exists t \ge N^2 : \frac{2}{1 + L^0(t)/|x_N|} \ge \varepsilon/2\right)$$
$$\le \mathbb{P}\left(\inf_{t \ge N^2} \left(1 + L^0(t)/t \cdot N^2/|x_N|\right) \le 4/\varepsilon\right),$$

Theorem 6.1 shows that $L^0(t)/t \to \lambda - \mu$ almost surely as $t \to +\infty$, and $N^2/|x_N|$ goes to infinity as N goes to infinity by choice of x_N . Hence almost surely,

$$\lim_{N \to +\infty} \inf_{t \ge N^2} \left(1 + L^0(t)/t \cdot N^2/|x_N| \right) = +\infty,$$

and the theorem is proved.

We now fill in the gap in this proof by proving Lemma 6.2.

PROOF OF LEMMA 6.2. Process X admits the following representation as the solution of a system of integral equations:

$$\begin{aligned} X_i(t) &= X_i(0) + \mathcal{N}_{\lambda_i}(t) - \int_0^t \mathbbm{1}_{\{X_i(s^-) \ge 1\}} \mathcal{N}_{\mu_i}(ds) \\ &+ \sum_{j \ne i} \int_0^t \sum_{k=1}^{X_j(s^-)} \mathcal{N}_{q_{ji}}^k(ds) - \sum_{j \ne i} \int_0^t \sum_{k=1}^{X_i(s^-)} \mathcal{N}_{q_{ij}}^k(ds) \quad 1 \le i \le n, \end{aligned}$$

where \mathcal{N}_{λ_i} and \mathcal{N}_{μ_i} , for i = 1, ..., n, are Poisson processes with respective parameters λ_i and μ_i , and for $(i, j) \in \{1, ..., n\}^2$, $i \neq j$, $(\mathcal{N}_{q_{ij}}^k, k \geq 1)$ is a sequence of Poisson processes with parameter q_{ij} , all these processes being independent.

Now using the same Poisson processes for X^x and X^y , it is easy to check that the inequalities $X_i^x(t) \ge X_i^y(t)$ true at t = 0 are preserved at each jump of any of the Poisson processes involved, and that $|X^x| - |X^y|$ is decreasing over time.

The previous results make it possible to establish the fluid regime of the system by studying the rescaled process \overline{X}_N defined by

(10)
$$\overline{X}_N(t) = \frac{X(Nt)}{N}, \quad t \ge 0.$$

In the following, \overline{L}_N denotes the rescaled number of particles, i.e., $\overline{L}_N(t) = L(Nt)/N$, and $\overline{\chi}_N = \overline{X}_N/\overline{L}_N$ is the corresponding proportions. Note that any fluid limit is discontinuous at 0+ (so that strictly speaking, X does not have any fluid limit), because Proposition 5.1 will show that the fluid limit is at π at time 0+, and Theorem 6.2 will imply that it stays forever proportional to π .

COROLLARY 6.1. Assume $\lambda > \mu$, and let $x : [0, +\infty[\rightarrow \mathbb{R}^n \text{ be defined by}]$

$$x(t) = (1 + (\lambda - \mu)t)\pi.$$

Then, for any sequence $(x_N, N \ge 1)$ with $|x_N| = N$, any s,t such that 0 < s < t and any $\varepsilon > 0$:

$$\lim_{N \to +\infty} \mathbb{P}_{x_N} \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x(u) \right\| \ge \varepsilon \right) = 0.$$

PROOF. Since the size of the initial state goes to infinity, Proposition 5.1 shows that for any $\delta > 0$, the event $\{T_{\delta} \leq t_{\delta}\}$ occurs with high probability. Since T_{δ} is a stopping time, the strong Markov property makes it possible to use $X_{T_{\delta}}$ as a new initial point, which is as close to equilibrium as desired. Since moreover the total number of customers did not significantly evolve in this time interval, this initial point will satisfy the hypotheses of Theorem 6.2, which makes it possible to conclude. Denote $\Delta_N(s,t)$ the distance of interest:

(11)
$$\Delta_N(s,t) = \sup_{s \le u \le t} \left\| \overline{X}_N(u) - x(u) \right\|.$$

First, the following decomposition makes it possible to consider all further convergences on the set $\{T_{\delta} \leq t_{\delta}\}$:

$$\mathbb{P}_{x_N}\left(\Delta_N(s,t) \ge \varepsilon\right) \le \mathbb{P}_{x_N}\left(\Delta_N(s,t) \ge \varepsilon, T_\delta \le t_\delta\right) + \mathbb{P}_{x_N}(T_\delta > t_\delta),$$

and the last term goes to 0 by Proposition 5.1. The strong Markov property used with the stopping time T_{δ} then shows that

$$\mathbb{P}_{x_N}\left(\Delta_N(s,t) \ge \varepsilon, T_\delta \le t_\delta\right) \le \mathbb{E}_{x_N}\left[\mathbb{P}_{X(T_\delta)}\left(\Delta_N(0,t) \ge \varepsilon\right)\right]$$

Now, we isolate the event of interest $\{|L(T_{\delta}) - |x_N|| \le \sqrt{N}\}$ by writing:

$$\mathbb{E}_{x_N} \left[\mathbb{P}_{X(T_{\delta})} \left(\Delta_N(0, t) \ge \varepsilon \right); \left| L(T_{\delta}) - |x_N| \right| \le \sqrt{N} \right] \\ \le \max_{\substack{y \in \mathbb{N}^n: \left| |y| - |x_N| \right| \le \sqrt{N} \\ \text{and } \|y/|y| - \pi\| \le \delta}} \mathbb{P}_y \left(\Delta_N(0, t) \ge \varepsilon \right),$$

therefore, if we note y_N the value that realizes this maximum (the set over which the maximum is considered is finite),

$$\mathbb{E}_{x_N}\left[\mathbb{P}_{X(T_{\delta})}\left(\Delta_N(0,t)\geq\varepsilon\right)\right]\leq\mathbb{P}_{x_N}\left(\left|L(T_{\delta})-|x_N|\right|\geq\sqrt{N}\right)+\mathbb{P}_{y_N}\left(\Delta_N(0,t)\geq\varepsilon\right).$$

The following inequality holds for any time $u \ge 0$ and any initial state:

$$|L(u) - L(0)| \le \mathcal{N}_{\lambda}(u) + \mathcal{N}_{\mu}(u) \stackrel{\text{def.}}{=} \mathcal{N}_{\lambda+\mu}(u),$$

and yields

$$\begin{aligned} \mathbb{P}_{x_N}\left(\left|L(T_{\delta}) - |x_N|\right| \geq \sqrt{N}\right) &\leq \mathbb{P}_{x_N}\left(\mathcal{N}_{\lambda+\mu}(T_{\delta}) \geq \sqrt{N}, T_{\delta} \leq t_{\delta}\right) + \mathbb{P}_{x_N}\left(T_{\delta} > t_{\delta}\right) \\ &\leq \mathbb{P}\left(\mathcal{N}_{\lambda+\mu}(t_{\delta}) \geq \sqrt{N}\right) + \mathbb{P}_{x_N}\left(T_{\delta} > t_{\delta}\right). \end{aligned}$$

This last sum vanishes, so that all is left to prove is that as $N \to +\infty$,

$$\mathbb{P}_{y_N}\left(\Delta_N(0,t) \ge \varepsilon\right) = \mathbb{P}_{y_N}\left(\sup_{0 \le u \le t} \left\|\overline{X}_N(u) - x(u)\right\| \ge \varepsilon\right) \to 0.$$

Note that the initial state y_N is now such that $|y_N|/N$ goes to 1 (because $|x_N| = N$ and $||y_N| - |x_N|| \le \sqrt{N}$), and $H(y_N/|y_N|, \pi)$ is as small as needed to apply Theorem 6.2, since $||y_N/|y_N| - \pi|| \le \delta$ and $\delta > 0$ is arbitrary small.

The triangular inequality and the definition of x give for any $0 \le u \le t$

$$\begin{aligned} \left\|\overline{X}_{N}(u) - x(u)\right\| &\leq \left\|\overline{X}_{N}(u) - \overline{L}_{N}(u)\pi\right\| + \left\|\left[\overline{L}_{N}(u) - (1 + (\lambda - \mu)u)\right]\pi\right\| \\ &\leq \left\|\overline{\chi}_{N}(u) - \pi\right\| \sup_{0 \leq u \leq t} \overline{L}_{N}(u) + \left\|\pi\right\| \sup_{0 \leq u \leq t} \left|\overline{L}_{N}(u) - (1 + (\lambda - \mu)u)\right|, \end{aligned}$$

and so

(12)
$$\mathbb{P}_{y_N} \left(\Delta_N(0,t) \ge \varepsilon \right) \le \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \| \overline{\chi}_N(u) - \pi \| \ge \varepsilon / \left(2 \sup_{0 \le u \le t} \overline{L}_N(u) \right) \right)$$
$$+ \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \left| \overline{L}_N(u) - (1 + (\lambda - \mu)u) \right| \ge \varepsilon / (2\|\pi\|) \right).$$

Under \mathbb{P}_{y_N} , a trivial upper bound for $\overline{L}_N(u)$ for $0 \le u \le t$ is given by

$$\overline{L}_N(u) \le \frac{1}{N} \left(|y_N| + \mathcal{N}_{\lambda}(Nt) \right) \stackrel{\text{def.}}{=} A_N(t),$$

therefore for any constant C > 0,

$$\mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t}\|\overline{\chi}_N(u)-\pi\|\geq \varepsilon/\left(2\sup_{0\leq u\leq t}\overline{L}_N(u)\right)\right)$$
$$\leq \mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t}\|\overline{\chi}_N(u)-\pi\|\geq \varepsilon/(2C)\right) + \mathbb{P}(A_N(t)\geq C).$$

For any $t \ge 0$, $A_N(t)$ converges almost surely to $1 + \lambda t$ as N goes to infinity, therefore $\mathbb{P}(A_N(t) \ge C)$ goes to 0 for $C = 2(1 + \lambda t)$. The other term vanishes as well. Indeed,

$$\mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t}\|\overline{\chi}_N(u)-\pi\|\geq \varepsilon/(2C)\right)=\mathbb{P}_{y_N}\left(T^{\varepsilon/(2C)}\leq Nt\right),$$

and by Lemma 2.1, there exists some $\varepsilon' > 0$ such that $T^{\varepsilon/(2C)} \ge T_H^{\varepsilon'}$, hence

$$\mathbb{P}_{y_N}\left(T^{\varepsilon/(2C)} \le Nt\right) \le \mathbb{P}_{y_N}\left(T_H^{\varepsilon'} \le Nt\right) \le \mathbb{P}_{y_N}\left(T_H^{\varepsilon'} < +\infty\right)$$

One can moreover assume $\varepsilon' < \varepsilon_0$ without loss of generality. Observe that so far, δ is arbitrary: it can be chosen small enough, say $\delta \leq \delta_0$ so that using Lemma 2.1, $H(y_N/|y_N|,\pi) \leq \varepsilon'/2$. Thus y_N satisfies the hypotheses of Theorem 6.2, which shows that $\mathbb{P}_{y_N}(T_H^{\varepsilon'} < +\infty)$, and hence the first term in the upper bound of (12), vanishes in the limit $N \to +\infty$.

The second term of (12) is easier to deal with. We reduce the problem to the event $\{\mathcal{T}_0 = +\infty\}$ by using the following upper bound:

(13)
$$\mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \left| \overline{L}(u) - (1 + (\lambda - \mu)u) \right| \ge \varepsilon / (2\|\pi\|) \right) \le \mathbb{P}_{y_N}(\mathcal{T}_0 < +\infty)$$
$$+ \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \left| \overline{L}(u) - (1 + (\lambda - \mu)u) \right| \ge \varepsilon / (2\|\pi\|), \mathcal{T}_0 = +\infty \right).$$

The first term $\mathbb{P}_{y_N}(\mathcal{T}_0 < +\infty)$ in the right hand side of (13) goes to 0 since, by Lemma 2.2, $\mathbb{P}_{y_N}(\mathcal{T}_0 < +\infty) \leq \mathbb{P}_{y_N}(T_H^{\varepsilon'} < +\infty)$ which has just been proved to vanish as $N \to +\infty$.

Because $L(u) = L(0) + \mathcal{N}_{\lambda}(u) - \mathcal{N}_{\mu}(u)$ for all $u \ge 0$ on $\{\mathcal{T}_0 = +\infty\}$, we get the following upper bound for the second term:

$$\mathbb{P}\left(\sup_{0\leq u\leq t}\left|\frac{1}{N}\left(|y_N|+\mathcal{N}_{\lambda}(Nu)-\mathcal{N}_{\mu}(Nu)\right)-(1+(\lambda-\mu)u)\right|\geq \varepsilon/(2\|\pi\|)\right),$$

and this term goes to 0 thanks to Doob's inequality. The proof is complete.

7. Stability of the Subcritical Regime

In this section we consider the subcritical regime $\lambda < \mu$ that is, the case when the input rate is smaller than the maximal output rate. As in the previous section, the key ingredients are the short and long term "homogenization" property and the M/M/1-like behavior of the total number of customers. The next lemma will be useful for establishing the fluid behavior of the system. It gives a control on the stopping time T_H^{ε} , or equivalently T^{ε} : with high probability, T_H^{ε} is larger than the time needed for a stable M/M/1 queue to empty.

LEMMA 7.1. Assume $\lambda < \mu$. Fix some a > 0 and let $(x_N, N \ge 1)$ be any sequence in \mathbb{N}^n such that

$$\lim_{N \to +\infty} \frac{|x_N|}{N} = a \text{ and } \lim_{N \to +\infty} H(x_N/|x_N|, \pi) = 0.$$

Then, for any $t < a/(\mu - \lambda)$ and any $\varepsilon < \varepsilon_0$, where ε_0 is given by Lemma 2.2,

$$\lim_{N \to +\infty} \mathbb{P}_{x_N} \left(T_H^{\varepsilon} \le Nt \right) = 0.$$

PROOF. Denote $H_N = H(x_N/|x_N|, \pi)$, and let $(\ell_N, N \ge 1)$ be a sequence of integers such that $N \gg \ell_N \gg NH_N$ and $\ell_N \gg \log N$ (such a sequence clearly exists, e.g., $\ell_N = N\sqrt{H_N} \vee (\log N)^2$). Proposition 3.2 with $\alpha = 1/N^2$ and $\delta = \varepsilon/2$ gives

(14)
$$\mathbb{E}_{x_N}\left(e^{-\theta T_H^{\varepsilon}/N^2}; L(T_H^{\varepsilon}) \ge \ell_N\right) \le C_{\varepsilon/2} e^{2n\log N + |x_N|H_N - \varepsilon \ell_N/2},$$

where the last bound goes to 0 by choice of ℓ_N . Let now τ_N be defined by $\tau_N = \inf\{t \ge 0 : L(t) \le \ell_N\}$. Since ℓ_N is an integer and L has jumps ± 1 , we have $L(\tau_N) = \ell_N$, and consequently, for any t > 0,

(15)
$$\mathbb{E}_{x_N}\left(e^{-\theta T_H^{\varepsilon}/N^2}; L(T_H^{\varepsilon}) \ge \ell_N\right) \ge \mathbb{E}_{x_N}\left(e^{-\theta T_H^{\varepsilon}/N^2}; T_H^{\varepsilon} \le \tau_N\right)$$

 $\ge e^{-\theta t/N} \mathbb{P}_{x_N}\left(T_H^{\varepsilon} \le \tau_N \land Nt\right).$

Inequalities (14) and (15) together imply that \mathbb{P}_{x_N} $(T_H^{\varepsilon} \leq \tau_N \wedge Nt)$ goes to 0 as N goes to infinity. Since

$$\mathbb{P}_{x_N} \left(T_H^{\varepsilon} \le Nt \right) \le \mathbb{P}_{x_N} \left(T_H^{\varepsilon} \le \tau_N \wedge Nt \right) + \mathbb{P}_{x_N} \left(\tau_N < Nt \right) + \mathbb{P}_{x_N} \left($$

all is left to prove is that $\mathbb{P}_{x_N}(\tau_N < Nt)$ goes to 0 if $t < a/(\mu - \lambda)$. Using the lower bound $L(t) \ge L(0) + \mathcal{N}_{\lambda}(t) - \mathcal{N}_{\mu}(t)$ from Proposition 4.1 and the fact that $|x_N| \ge (\mu - \lambda)Nt + \ell_N$ for N large enough if $t < a/(\mu - \lambda)$, we get for such a t

$$\mathbb{P}_{x_N} \left(\tau_N < Nt \right) \le \mathbb{P}_{x_N} \left(\exists s \in [0, Nt] : L(0) + \mathcal{N}_{\lambda}(s) - \mathcal{N}_{\mu}(s) \le \ell_N \right)$$
$$\le \mathbb{P}_{x_N} \left(\sup_{0 \le s \le Nt} \left(\mathcal{N}_{\mu}(s) - \mathcal{N}_{\lambda}(s) - (\mu - \lambda)s \right) \ge |x_N| - (\mu - \lambda)Nt - \ell_N \right)$$
$$\le \mathbb{P}_{x_N} \left(\sup_{0 \le s \le Nt} \left(\mathcal{N}_{\mu}(s) - \mathcal{N}_{\lambda}(s) - (\mu - \lambda)s \right)^2 \ge (|x_N| - (\mu - \lambda)Nt - \ell_N)^2 \right).$$

Since $(\mathcal{N}_{\mu}(s) - \mathcal{N}_{\lambda}(s) - (\mu - \lambda)s, s \ge 0)$ is a martingale, Doob's inequality yields that the last term is in turn upper bounded by

$$\frac{\operatorname{\mathbb{V}ar}\left(\mathcal{N}_{\lambda}(Nt)\right) + \operatorname{\mathbb{V}ar}\left(\mathcal{N}_{\mu}(Nt)\right)}{(|x_{N}| - (\mu - \lambda)Nt - \ell_{N})^{2}} \sim \frac{(\lambda + \mu)Nt}{(a - (\mu - \lambda)t)^{2}N^{2}} \to 0,$$

which completes the proof.

The fluid behavior can now be established. Recall that the rescaled process $(\overline{X}_N(t))$ is defined by $\overline{X}_N(t) = X(Nt)/N$ for any $t \ge 0$. In what follows, for $u \in \mathbb{R}$, u^+ denotes $\max(u, 0)$.

PROPOSITION 7.1. Let $x: [0, +\infty] \to \mathbb{R}^n$ be defined by

$$x(t) = \left(1 + (\lambda - \mu)t\right)^+ \pi.$$

Then for all 0 < s < t and all $\varepsilon > 0$:

$$\lim_{N \to +\infty} \left[\max_{x \in \mathbb{N}^n : |x| = N} \mathbb{P}_x \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x(u) \right\| \ge \varepsilon \right) \right] = 0.$$

PROOF. Lemma 7.1 makes it possible to study the system for $t < 1/(\mu - \lambda)$. An additional coupling argument, involving larger initial states, is then required to show that fluid limits stay at 0 after that time. For this technical reason, initial states of size equivalent to aN for some a > 0 will be considered, and the following more general result will be established: For a > 0, let $x_a : [0, +\infty[\rightarrow \mathbb{R}^n]$ be defined by

$$x_a(t) = \left(a + (\lambda - \mu)t\right)^+ \pi.$$

It will be proved that for any a > 0, any s, t with 0 < s < t and all $\varepsilon > 0$:

$$\lim_{N \to +\infty} \left[\max_{x \in \mathbb{N}^n : |x| = \lfloor aN \rfloor} \mathbb{P}_x \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x_a(u) \right\| \ge \varepsilon \right) \right] = 0,$$

where the notations of the previous section are used.

First assume $t < t_a = a/(\mu - \lambda)$, and set $\Delta_N(s, t) = \sup_{s \le u \le t} ||\overline{X}_N(u) - x_a(u)||$: the first steps of the proof are similar to the underloaded regime, namely using the strong Markov property to replace the arbitrary initial state by some initial state with low entropy. More precisely, let δ_N and t_N be as in Proposition 5.2: for any $x \in \mathbb{N}^n$ with |x| = N, one has

$$\mathbb{P}_x\left(\Delta_N(s,t) \ge \varepsilon\right) \le \mathbb{P}_x\left(\Delta_N(s,t) \ge \varepsilon, T_{\delta_N} \le Ns\right) + \mathbb{P}_x(T_{\delta_N} > Ns).$$

Since t_N/N goes to 0, Proposition 5.2 gives that the last term $\mathbb{P}_x(T_{\delta_N} > Ns)$ goes to 0 uniformly in $x \in \mathbb{N}^n$ with $|x| = \lfloor aN \rfloor$. As for the first term, we write

$$\mathbb{P}_{x}\left(\Delta_{N}(s,t) \geq \varepsilon, T_{\delta_{N}} \leq Ns\right) \leq \mathbb{E}_{x}\left[\mathbb{P}_{X(T_{\delta_{N}})}\left(\Delta_{N}(0,t) \geq \varepsilon\right)\right]$$
$$\leq \mathbb{P}_{y_{N}}\left(\Delta_{N}(0,t) \geq \varepsilon\right) + \mathbb{P}_{x}\left(\left|L(T_{\delta_{N}}) - L(0)\right| \geq \sqrt{N}\right),$$

where $y_N \in \mathbb{N}^n$ is such that

$$\mathbb{P}_{y_N}\left(\Delta_N(0,t) \ge \varepsilon\right) = \max_{\substack{y \in \mathbb{N}^n : \left| |y| - \lfloor aN \rfloor \right| \le \sqrt{N} \\ \text{and } \|y/|y| - \pi \| \le \delta_N}} \mathbb{P}_y\left(\Delta_N(0,t) \ge \varepsilon\right).$$

Because $T_{\delta_N} \leq t_N$ with high probability, and because $t_N/\sqrt{N} \to 0$, one can show similarly as in Section 6 that as N goes to infinity,

$$\max_{x \in \mathbb{N}^n : |x| = \lfloor aN \rfloor} \mathbb{P}_x \left(\left| L(T_{\delta_N}) - L(0) \right| \ge \sqrt{N} \right) \to 0.$$

Along the same lines as in the overloaded case, one gets, by introducing the term $\overline{L}_N(u)\pi$ that for any C>0

(16)
$$\mathbb{P}_{y_N} \left(\Delta_N(0,t) \ge \varepsilon \right) \le \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \| \overline{\chi}_N(u) - \pi \| \ge \varepsilon / (2C) \right)$$
$$+ \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \overline{L}_N(u) \ge C \right) + \mathbb{P}_{y_N} \left(\sup_{0 \le u \le t} \left| \overline{L}_N(u) - (a + (\lambda - \mu)u) \right| \ge \varepsilon / (2\|\pi\|) \right).$$

Note that since $t < t_a$, $x_a(u) = (a + (\lambda - \mu)u)\pi$ for $0 \le u \le t$. For C large enough, $\mathbb{P}_{y_N}(\sup_{0\le u\le t} \overline{L}_N(u) \ge C)$ goes to 0 as N goes to infinity. Moreover, Lemma 2.1 gives

$$\mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t}\|\overline{\chi}_N(u)-\pi\|\geq \varepsilon/\left(2C\right)\right) = \mathbb{P}_{y_N}\left(T^{\varepsilon/(2C)}\leq Nt\right)\leq \mathbb{P}_{y_N}\left(T_H^{\varepsilon'}\leq Nt\right)$$

for some $\varepsilon' > 0$ that can be assumed to satisfy $\varepsilon' < \varepsilon_0$. Since the sequence $(y_N, N \ge 1)$ satisfies the hypotheses of Lemma 7.1 and since $t < t_a$, this last upper bound goes to 0. Moreover, since

$$\mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t} \left|\overline{L}_N(u) - (a + (\lambda - \mu)u)\right| \geq \varepsilon/(2\|\pi\|)\right) \leq \mathbb{P}_{y_N}(\mathcal{T}_0 \leq Nt) \\ + \mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t} \left|\overline{L}_N(u) - (a + (\lambda - \mu)u)\right| \geq \varepsilon/(2\|\pi\|), \mathcal{T}_0 > Nt\right),$$

we conclude, using Lemma 7.1 together with Lemma 2.2 for the first term and Doob's inequality for the second one, that

$$\mathbb{P}_{y_N}\left(\sup_{0\leq u\leq t}\left|\overline{L}_N(u)-(a+(\lambda-\mu)u)\right|\geq \varepsilon/(2\|\pi\|)\right)\to 0.$$

The proof in the case $0 < s < t < t_a$ is thus complete.

To conclude in the other cases, a monotonicity argument derived from the above Lemma 6.2 is used. Let 0 < s < t and $t \ge t_a$, and assume in a first step that $t - s < \varepsilon/(2(\mu - \lambda))$. In addition, let $b > (\mu - \lambda)t$ be fixed, and let $t_b = b/(\mu - \lambda)$ be the corresponding time. Note that $t \ge t_a$ implies that b > a, so that for any $x \in \mathbb{N}^n$ with $|x| = \lfloor aN \rfloor$, there exists some $y \in \mathbb{N}^n$ such that $y \ge x$ componentwise and $|y| = \lfloor bN \rfloor$. For such x, y, Lemma 6.2 shows that X^x and X^y can be coupled in such a way that $|X^x(t)| \le |X^y(t)|$ for any $t \ge 0$. Hence for any $u \ge s$, using the inequality $||v|| \le |v| \le n ||v||$ for any $v \in \mathbb{R}^n$, one gets

$$\begin{aligned} \left\|\overline{X}_{N}^{x}(u) - x_{a}(u)\right\| &\leq \left|\overline{X}_{N}^{x}(u)\right| + \left|x_{a}(s)\right| \leq \left|\overline{X}_{N}^{y}(u)\right| + \left|x_{a}(s)\right| \\ &\leq n \left\|\overline{X}_{N}^{y}(u) - x_{b}(u)\right\| + \left|x_{b}(s)\right| + \left|x_{a}(s)\right|. \end{aligned}$$

By definition

$$|x_b(s)| + |x_a(s)| = (\mu - \lambda)(t_b - s) + (\mu - \lambda)(t_a - s)^+ \le 2(\mu - \lambda)(t_b - s).$$

This yields in turn

$$\mathbb{P}_x\left(\sup_{s\leq u\leq t}\left\|\overline{X}_N(u)-x_a(u)\right\|\geq\varepsilon\right)\leq\mathbb{P}_y\left(\sup_{s\leq u\leq t}\left\|\overline{X}_N(u)-x_b(u)\right\|\geq\varepsilon''\right)$$

where $\varepsilon'' = (\varepsilon - 2(\mu - \lambda)(t_b - s))/n$, and finally

$$\max_{x \in \mathbb{N}^n : |x| = \lfloor aN \rfloor} \mathbb{P}_x \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x_a(u) \right\| \ge \varepsilon \right) \\ \le \max_{y \in \mathbb{N}^n : |y| = \lfloor bN \rfloor} \mathbb{P}_y \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x_b(u) \right\| \ge \varepsilon'' \right).$$

Since it has been assumed that $t - s < \varepsilon/(2(\mu - \lambda))$, $b > (\mu - \lambda)t$ can be chosen small enough so that $\varepsilon'' > 0$. Since $t < t_b$, the first part of the proof implies that

$$\lim_{N \to +\infty} \left[\max_{y \in \mathbb{N}^n : |y| = \lfloor bN \rfloor} \mathbb{P}_y \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - x_b(u) \right\| \ge \varepsilon'' \right) \right] = 0.$$

This proves in particular that when 0 < s < t and $t - s < \varepsilon/(2(\mu - \lambda)))$, then $\max_{|x|=\lfloor aN \rfloor} \mathbb{P}_x (\Delta_N(s,t) \ge \varepsilon) \to 0$. It is now left to extend this result to any s, t such that s < t, which is a consequence of the following decomposition:

$$\max_{x \in \mathbb{N}^n : |x| = \lfloor aN \rfloor} \mathbb{P}_x \left(\Delta_N(s, t) \ge \varepsilon \right) \le \sum_{j=1}^q \left(\max_{x \in \mathbb{N}^n : |x| = \lfloor aN \rfloor} \mathbb{P}_x \left(\Delta_N(s_{j-1}, s_j) \ge \varepsilon \right) \right)$$

where $s_0 = s < s_1 < \ldots < s_q = t$ and $s_j - s_{j-1} < \varepsilon/(2(\mu - \lambda))$ for $1 \le j \le q$. Indeed, it has just been shown that each term of this finite sum goes to 0.

REMARK 7.1. It can be proved that in the critical case $\lambda = \mu > 0$, the fluid limit is constant and equal to π , i.e., if $\lambda = \mu > 0$, then for all 0 < s < t and all $\varepsilon > 0$,

$$\lim_{N \to +\infty} \left[\max_{x \in \mathbb{N}^n : |x| = N} \mathbb{P}_x \left(\sup_{s \le u \le t} \left\| \overline{X}_N(u) - \pi \right\| \ge \varepsilon \right) \right] = 0$$

This convergence follows readily from Proposition 7.1 and the following coupling. For $0 \leq \eta \leq \lambda$, if X^{η} is a subcritical process with arrival rate $\lambda - \eta$ and departure rate $\lambda = \mu$, then X and X^{η} can be coupled in such a way that $\|\overline{X}(t) - \overline{X}^{\eta}(t)\| \leq N_{\eta}(t)$ for all $t \geq 0$, where N_{η} is a Poisson process with intensity η .

Note that the behavior of the fluid limit in the critical case does not make it possible to infer the stability or the transience of process X. The analogy with the M/M/1 queue nevertheless suggests that it could be recurrent null in this case.

In contrast, the behavior of the fluid limit shows that X is ergodic in the subcritical case $\lambda < \mu$:

PROPOSITION 7.2. When $\lambda < \mu$, the Markov process X is ergodic.

PROOF. According to [**Rob03**, Corollary 9.8 p. 259], it is enough to show that for some deterministic time T > 0,

$$\lim_{N \to +\infty} \max_{x \in \mathbb{N}^n : |x| = N} \mathbb{E}_x(\overline{L}_N(T)) = 0.$$

Recall that $\overline{L}_N(T) = L(NT)/N$, and let $\varepsilon > 0$ be fixed: then, for $x \in \mathbb{N}^n$ with |x| = N:

$$\mathbb{E}_x\left(\overline{L}_N(T)\right) \le \varepsilon + \mathbb{E}_x\left(\overline{L}_N(T); \overline{L}_N(T) > \varepsilon\right) \le \varepsilon + (1 + \lambda T)\mathbb{P}_x\left(\overline{L}_N(T) > \varepsilon\right) \\ + \frac{1}{N}\mathbb{E}_x\left(\mathcal{N}_\lambda(NT) - \lambda NT; \overline{L}_N(T) > \varepsilon\right).$$

where the second inequality comes from $\overline{L}_N(T) \leq \overline{L}_N(0) + \mathcal{N}_\lambda(NT)/N$. For any $T \geq 0$, using Cauchy-Schwartz inequality, an upper bound on the last term is given by

$$\frac{1}{N}\mathbb{E}_x\left(\mathcal{N}_\lambda(NT) - \lambda NT; \overline{L}_N(T) > \varepsilon\right) \le \frac{1}{N}\mathbb{E}\left(|\mathcal{N}_\lambda(NT) - \lambda NT|\right) \le \sqrt{\frac{\lambda T}{N}},$$

so that finally

$$\max_{x \in \mathbb{N}^n : |x|=N} \mathbb{E}_x \left(\overline{L}_N(T) \right) \le \varepsilon + (1+\lambda T) \max_{x \in \mathbb{N}^n : |x|=N} \mathbb{P}_x \left(\overline{L}_N(T) > \varepsilon \right) + \sqrt{\frac{\lambda T}{N}},$$

and all is left to prove is that for some T > 0, $\max_{|x|=N} \mathbb{P}_x(\overline{L}_N(T) > \varepsilon)$ goes to 0 as N grows to infinity: this is a direct consequence of Proposition 7.1 with $T = 1/(\mu - \lambda)$ since x(T) = 0. The proof is now complete.

Appendix A. Martingale Construction

This appendix is devoted to proving Theorem 3.1, which states the existence of a fundamental family of local martingales. In Proposition A.1, we first establish the harmonicity of a special function g, which has an integral form. Then a change of variables leads to the local martingale introduced in Theorem 3.1.

A.1. An Integral Harmonic Function. The starting point is the generator Ω of the Markov process X given, for any $x \in \mathbb{N}^n$ and any function $f : \mathbb{R}^n \to \mathbb{R}$, by:

$$\Omega(f)(x) = \sum_{i=1}^{n} \lambda_i \big(f(x+e_i) - f(x) \big) + \sum_{i=1}^{n} \mu_i \big(f(x-e_i) - f(x) \big) \mathbb{1}_{\{x_i > 0\}} \\ + \sum_{1 \le i \ne j \le n} q_{ij} x_i \big(f(x+e_j - e_i) - f(x) \big).$$

In addition to the irreducibility of $Q = (q_{ij})_{1 \le i,j \le n}$, we will require that Q is diagonalizable in \mathbb{C} , i.e., that there exists a set $(\omega_j, 1 \le j \le n)$ of eigenvectors of Q that generate \mathbb{R}^n . The complex square matrix $\omega = (\omega_{i,j})_{1 \le i,j \le n}$ where $\omega_j = (\omega_{i,j})_{1 \le i \le n}$ is invertible.

We can assume without loss of generality that $\omega_n = 1$, denoting by 1 the vector in \mathbb{R}^n with all coordinates equal to 1, so that ω_n is associated to the null eigenvalue; more generally, for $1 \leq j \leq n$, θ_j will denote the (possibly complex) eigenvalue associated to ω_j . The negative trace of Q is then given by $-\theta = \sum_{1}^{n} \theta_i$ with $\theta > 0$.

In the sequel \mathcal{H} will denote the hyperplane of \mathbb{R}^n defined by

$$\mathcal{H} = \left\{ v \in \mathbb{R}^n : \sum_{i=1}^n \pi_i v_i = 0 \right\}.$$

For j = 1, ..., n - 1, $\omega_j \in \mathcal{H}$ since $Q\omega_j = \theta_j\omega_j$ for $\theta_j \neq 0$ implies (in a matricial form, where π is a row and ω_j a column): $\pi\omega_j = (\theta_j)^{-1}\pi Q\omega_j$ which is 0 since $\pi Q = 0$. These n - 1 eigenvectors then generate \mathcal{H} .

We recall some notations and results of Section 3. $(P_t, t \in \mathbb{R})$ denotes the Q-generated Markov semi-group of linear operators in \mathbb{R}^n : $P_t = e^{tQ}$, extended to all real indices t into a group. Each P_t has eigenvalues $e^{\theta_j t}$ and eigenvectors ω_j , $j = 1, \ldots, n$. For any $v \in \mathbb{R}^n$ and $t \ge 0$, we define

$$\phi(v,t) = (\phi_i(v,t), 1 \le i \le n) = P_{-t}v.$$

If $v \in \mathbb{R}^n$ and $\varphi(v, \cdot)$ is any primitive of $\sum_{i=1}^n [\mu_i \phi_i(v, \cdot)/(1 + \phi_i(v, \cdot)) - \lambda_i \phi_i(v, \cdot)]$ on some open subset V of $\{t \geq 0 : \forall i = 1, ..., n, 1 + \phi_i(v, t) \neq 0\}$, then the function $h_v(t, x)$ defined by

(17)
$$h_v(t,x) = e^{\varphi(v,t)} \prod_{i=1}^n \left(1 + \phi_i(v,t)\right)^{x_i}$$

is space-time harmonic with respect to Ω in the domain $V \times \mathbb{N}^{*n}$ (see Proposition 3.1).

The suitable domain of integration for constructing our martingale will be:

$$\mathcal{D}(t) = \{ v \in \mathcal{H} : \mathbb{1} + \phi(v, t) > 0 \}, \quad t \in \mathbb{R},$$

where, for any $u \in \mathbb{R}^n$, $u \ge 0$ (resp. u > 0) means that $u_i \ge 0$ (resp. $u_i > 0$) for every i = 1, ..., n.

For each $t \in \mathbb{R}$, $\mathcal{D}(t)$ is an open subset of \mathcal{H} . Moreover, it is clear from the definition of $\mathcal{D}(t)$ and from the invariance of \mathcal{H} under the group of operators $(P_s, s \in \mathbb{R})$ that, for any $v \in \mathbb{R}^n$ and any $t \ge 0$,

$$v \in \mathcal{D}(t) \Longleftrightarrow P_{-t}v \in \mathcal{D}(0).$$

So, for any $t \in \mathbb{R}$, $\mathcal{D}(t) = P_t(\mathcal{D}(0))$. Then, since $\mathcal{D}(0) = \{v \in \mathbb{R}^n : \sum_{i=1}^n \pi_i v_i = 0, 1+v > 0\}$ is clearly bounded, each $\mathcal{D}(t) = P_t(\mathcal{D}(0))$ for $t \in \mathbb{R}$ is bounded as well.

Define the subset \mathcal{A} of $\mathcal{H} \times \mathbb{R}$ by:

$$\mathcal{A} = \{ (v, t) : t \in \mathbb{R} \text{ and } v \in \mathcal{D}(t) \}.$$

The first step is to show that the following choice of φ makes sense:

(18)
$$\varphi(v,t) = \int_{-\infty}^{t} \sum_{i=1}^{n} \left(\mu_i \frac{\phi_i(v,s)}{1 + \phi_i(v,s)} - \lambda_i \phi_i(v,s) \right) ds$$

for $(v,t) \in \mathcal{A}$. This is the object of the following two lemmas, which will also give some regularity properties of φ in view of Proposition A.1.

LEMMA A.1. If $(v,t) \in \mathbb{R}^n \times \mathbb{R}$ satisfies $1 + \phi(v,t) \ge 0$ and $v \ne -1$, then $1 + \phi(v,s) > 0$ for all s < t. As a consequence:

(19)
$$t > s \Longrightarrow \mathcal{D}(t) \subset \mathcal{D}(s), \quad s, t \in \mathbb{R},$$

and

(20)
$$\mathcal{D}(t_0) = \bigcup_{t > t_0} \mathcal{D}(t), \quad t_0 \in \mathbb{R}.$$

PROOF. Let first remark that the irreducibility of Q implies that, for any r > 0 and any $(i, j) \in \{1, \ldots, n\}$, the probability $P_r(i, j)$ that a Markov process with generator Q initiated at i is in state j at time r is positive. Indeed, if $i = i_0, i_1, \ldots, i_k = j$ is a path from i to j such that $q_{i_{l-1},i_l} > 0$ for $l = 1, \ldots, k$, then there is a positive probability that the process has exactly followed this path by time r.

This implies that $P_r u > 0$ for any r > 0 and $u \in \mathbb{R}^n$ such that $u \ge 0$ and $u \ne 0$. Now let (v, t) satisfy the hypotheses in the lemma, then $1 + \phi(v, t) \ne 0$ since

$$0 \neq \mathbb{1} + v = P_t \big(\mathbb{1} + P_{-t} v \big) = P_t \big(\mathbb{1} + \phi(v, t) \big),$$

and the previous property applied to $u = 1 + \phi(v, t)$ and r = t - s for s < t gives,

$$1 + \phi(v, s) = 1 + P_{-s}v = P_{t-s}(1 + P_{-t}v) = P_{t-s}(1 + \phi(v, t)) > 0.$$

The implication $t > s \Longrightarrow \overline{\mathcal{D}(t)} \subset \mathcal{D}(s)$ follows, noticing that $-\mathbb{1} \notin \mathcal{H}$.

To show (20) for $t_0 \in \mathbb{R}$, note that $\mathcal{D}(t_0)$ contains the right hand side union by (19), and that the reverse holds since the inequality $1 + \phi(v, t_0) > 0$ extends to some neighborhood of t_0 for $v \in \mathcal{D}(t_0)$.

LEMMA A.2. (i) For any $i \in \{1, ..., n\}$, the two integrals

$$\int_{-\infty}^{0} \phi_i(v,s) ds \text{ and } \int_{-\infty}^{0} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds$$

are well defined for $v \in \mathcal{D}(0)$, continuous as functions of v on this domain and respectively bounded and bounded above on $\mathcal{D}(0)$.

The function φ_0 can then be defined on $\mathcal{D}(0)$ by

$$\varphi_0(v) = \int_{-\infty}^0 \sum_{i=1}^n \left(\mu_i \frac{\phi_i(v,s)}{1 + \phi_i(v,s)} - \lambda_i \phi_i(v,s) \right) ds$$

and is continuous and bounded above on $\mathcal{D}(0)$.

(ii) The function φ given by (18) is well defined for $(v, t) \in \mathcal{A}$ and satisfies:

$$\varphi(v,t) = \varphi_0(P_{-t}v) \qquad (v,t) \in \mathcal{A}.$$

The function φ is bounded above on \mathcal{A} and continuous with respect to $v \in \mathcal{D}(t)$ for fixed $t \in \mathbb{R}$.

PROOF. (i) Notice that, for fixed $v \in \mathbb{R}^n$, the map $s \mapsto \phi(v, s) = e^{-sQ}v$ is continuous on \mathbb{R} (with values in \mathbb{R}^n). Moreover, if $v \in \mathcal{H}$, it has a fast decay as s tends to $-\infty$ as a consequence of the exponential fast convergence of $P_t(i, \cdot)$ to π (already used in Section 5):

There exist some positive constants η and B_1 such that, for any $s \leq 0$

(21)
$$\max_{1 \le i,j \le n} |P_{-s}(i,j) - \pi_j| \le B_1 \cdot e^{\eta s}.$$

This gives, for $s \leq 0$ and $v \in \mathcal{H}$,

(22)
$$\|\phi(v,s)\| \le B_2 \cdot e^{\eta s} \|v\|,$$

where $B_2 = nB_1$, which ensures the existence of the vectorial integral $\int_{-\infty}^0 \phi(v, s) ds$ for any $v \in \mathcal{H}$. This integral is continuous with respect to v in \mathcal{H} since, for $v \in \mathcal{H}$,

$$\int_{-\infty}^{0} P_{-s} v ds = \int_{-\infty}^{0} \left(P_{-s} - \Pi \right) v ds = \left(\int_{-\infty}^{0} \left(P_{-s} - \Pi \right) ds \right) v ds$$

where Π is the square matrix with all lines equal to $\pi = (\pi_1, \ldots, \pi_n)$ and the last matricial integral has a coefficientwise meaning (and is well defined due to (21)). This shows the integral $\int_{-\infty}^{0} \phi(v, s) ds$ as a linear function of $v \in \mathcal{H}$, thus proving its continuity with respect to $v \in \mathcal{H}$. The boundedness of this function on $\mathcal{D}(0)$ follows since $\mathcal{D}(0)$ has compact closure in \mathcal{H} .

For the second integral, Lemma A.1 and the condition $v \in \mathcal{D}(0)$ first ensure that $1 + \phi(v, s) > 0$ for $s \leq 0$. The existence of this integral then again follows from the continuity of $s \mapsto \phi(v, s)$ and from the exponential decay in (22).

Let us now begin by proving that it is bounded above on $\mathcal{D}(0)$, writing

$$\int_{-\infty}^{0} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds = \int_{-\infty}^{-1} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds + \int_{-1}^{0} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds$$

and upperbounding each term.

It is easy for the second one, since $1 + \phi_i(v, s) > 0$ implies that $\phi_i(v, s)/(1 + \phi_i(v, s)) \leq 1$. In particular,

$$\int_{-1}^{0} \frac{\phi_i(v,s)}{1 + \phi_i(v,s)} ds \le 1.$$

The first term can be extended to $v \in \overline{\mathcal{D}(0)}$ (again by Lemma A.1 and by the exponential decay in (22)) and can be shown to be bounded on $\overline{\mathcal{D}(0)}$. Indeed, for

 $v \in \overline{\mathcal{D}(0)}$ and $s \leq -1$, $1 + \phi_i(v, s)$ is positive and tends to 1 as s tends to $-\infty$ uniformly in $v \in \overline{\mathcal{D}(0)}$ since, by (22),

(23)
$$\sup_{v\in\overline{\mathcal{D}(0)}} \|\phi(v,s)\| \le B_2 \cdot e^{\eta s} \sup_{v\in\overline{\mathcal{D}(0)}} \|v\|.$$

Then $1 + \phi_i(v, s)$ is bounded below by some positive δ for $(v, s) \in \mathcal{D}(0) \times]-\infty, -1]$ (by (23), this is the case on $\overline{\mathcal{D}(0)} \times]-\infty, -\kappa$] for κ large enough, and $\overline{\mathcal{D}(0)} \times [-\kappa, -1]$ is compact), and the following bound holds for $v \in \overline{\mathcal{D}(0)}$, using (23):

$$\left|\int_{-\infty}^{-1} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds\right| \leq \frac{B_2}{\delta} \cdot \sup_{u \in \overline{\mathcal{D}(0)}} \|u\| \int_{-\infty}^{-1} e^{\eta s} ds = \frac{B_2 e^{-\eta}}{\delta \eta} \sup_{u \in \overline{\mathcal{D}(0)}} \|u\|.$$

Let us now show the continuity of $\int_{-\infty}^{0} \phi_i(v,s)/(1+\phi_i(v,s))ds$ with respect to $v \in \mathcal{D}(0)$, using the continuity of $\phi_i(v,s)$ for fixed s, together with Lebesgue's dominated convergence theorem. The difficulty is that $\phi_i(v,s)/(1+\phi_i(v,s))$ is not clearly dominated uniformly in $v \in \mathcal{D}(0)$ by some integrable function of son $]-\infty, 0]$, since for s close to 0 and v close to the the portion of $\partial \mathcal{D}(0)$ where $1+v_i = 0$, the ratio $\phi_i(v,s)/(1+\phi_i(v,s))$ goes to infinity and is not easily controlled.

It is however possible to show local domination, using (20) in the particular case $t_0 = 0$ and dominating the integrand on each $\mathcal{D}(t), t > 0$. This will prove continuity on each $\mathcal{D}(t), t > 0$, hence continuity on $\mathcal{D}(0)$ since the $\mathcal{D}(t), t > 0$, are open subsets of $\mathcal{D}(0)$. The domination uses the same argument as in the last point: if t > 0, then

$$\delta(t) = \inf \left\{ 1 + \phi_i(v, s), \, (v, s) \in \overline{\mathcal{D}(t)} \times] - \infty, 0 \right\}$$

is positive. Then, for any $v \in \mathcal{D}(t)$ and $s \leq 0$,

$$\left|\frac{\phi_i(v,s)}{1+\phi_i(v,s)}\right| \leq \frac{B_2}{\delta(t)} \cdot e^{\eta s} \sup_{u \in \overline{\mathcal{D}(t)}} \|u\|$$

where the right hand side is integrable on $]-\infty, 0]$, and hence provides the required domination.

(ii) We use the group structure of $(P_s, s \in \mathbb{R})$ to rewrite both integrals as:

$$\int_{-\infty}^{t} \phi_i(v, s) ds = \int_{-\infty}^{0} \phi_i(v, s+t) ds = \int_{-\infty}^{0} \phi_i(P_{-t}v, s) ds$$

and

$$\int_{-\infty}^{t} \frac{\phi_i(v,s)}{1+\phi_i(v,s)} ds = \int_{-\infty}^{0} \frac{\phi_i(v,s+t)}{1+\phi_i(v,s+t)} ds = \int_{-\infty}^{0} \frac{\phi_i(P_{-t}v,s)}{1+\phi_i(P_{-t}v,s)} ds,$$

which ensures their existence for $P_t v \in \mathcal{D}(0)$, i.e., $v \in \mathcal{D}(t)$ or equivalently $(v, t) \in \mathcal{A}$, and proves the connexion between φ and φ_0 , hence the stated properties of φ .

The function φ is now legitimately defined by (18) for $t \in \mathbb{R}$ and $v \in \mathcal{D}(t)$. For fixed $t \in \mathbb{R}$, $\varphi(\cdot, t)$ is continuous (hence measurable) with respect to v in $\mathcal{D}(t)$ and bounded above on this domain.

Reversely, for fixed $v \in \mathcal{D}(0)$, $\varphi(v, \cdot)$ is clearly \mathcal{C}^1 on the interval $] - \infty, t_v[$, where $t_v = \sup\{t \ge 0 : \mathbb{1} + \phi(v, t) > 0\} > 0$ (by continuity of $\phi(v, \cdot)$ on \mathbb{R}), and $\partial \varphi(v, t) / \partial t = \sum_{i=1}^{n} [\mu_i \phi_i(v, t) / (1 + \phi_i(v, t)) - \lambda_i \phi_i(v, t)]$ on this interval. The following proposition constitutes the first step in defining a new spacetime harmonic function, obtained by integrating with respect to $v \in \mathcal{D}(t)$ the parametrized family of functions h_v given by (17), with φ given by (18).

As $\mathcal{D}(t) \subset \mathcal{H}$ and \mathcal{H} is an n-1 dimensional subspace of \mathbb{R}^n which is isomorphic to \mathbb{R}^{n-1} through the one to one linear mapping

$$H: \begin{array}{ccc} \mathbb{R}^{n-1} & \longrightarrow & \mathcal{H} \\ u & \longmapsto & \hat{u} = (u, -\sum_{i=1}^{n-1} \pi_i u_i / \pi_n) \end{array},$$

the new harmonic function will rather take the form of an integral over the following subset of \mathbb{R}^{n-1} :

$$\mathcal{C}(t) = H^{-1}(\mathcal{D}(t)) = \left\{ u \in \mathbb{R}^{n-1} : \hat{u} \in \mathcal{D}(t) \right\} = \left\{ u \in \mathbb{R}^{n-1} : \mathbb{1} + \phi(\hat{u}, t) > 0 \right\}.$$

PROPOSITION A.1. For any locally Lebesgue-integrable f on \mathbb{R}^{n-1} , the function g(t,x) given by the formula

(24)
$$g(t,x) = \int_{\mathcal{C}(t)} h_{\hat{u}}(t,x) f(u) du = \int_{\mathcal{C}(t)} e^{\varphi(\hat{u},t)} \cdot \prod_{i=1}^{n} (1 + \phi_i(\hat{u},t))^{x_i} \cdot f(u) du$$

is space-time harmonic in the domain $[0, +\infty[\times\mathbb{N}^{*n}]$.

PROOF. The function g is well-defined on $[0, +\infty[\times\mathbb{N}^{*n}]$. Indeed $\prod_{1}^{n} (1 + \phi_i(\hat{\cdot}, t))^{x_i}$ is continuous on \mathbb{R}^{n-1} for fixed $t \geq 0$ and $x \in \mathbb{N}^{*n}$, hence bounded on the bounded set $\mathcal{C}(t)$ (since $\mathcal{C}(t)$ corresponds to $\mathcal{D}(t)$ through H^{-1}). By Lemma A.2, $e^{\varphi(\hat{\cdot}, t)}$ also is continuous and bounded on $\mathcal{C}(t)$ since $\varphi(\cdot, t)$ is continuous and bounded on $\mathcal{O}(t)$. Then, since f is locally integrable, the product of these three functions is integrable on $\mathcal{C}(t)$.

We have to show that $\partial g/\partial t$ exists and satisfies $\partial g(t, x)/\partial t + \Omega(g(t, \cdot))(x) = 0$. As a rough argument, one expects that

(25)
$$\frac{\partial g}{\partial t}(t,x) = \int_{\mathcal{C}(t)} \frac{\partial h_{\hat{u}}}{\partial t}(t,x) \cdot f(u) du.$$

Indeed the additional derivation term resulting from the *t*-dependency of the domain C(t) is bound to vanish, since $h_{\hat{u}}(t,x)$ is zero for u on the frontier of C(t) (recall that $x \in \mathbb{N}^{*n}$). Therefore, since Ω commutes with integration,

$$\frac{\partial g}{\partial t}(t,x) + \Omega(g(t,\cdot))(x) = \int_{\mathcal{C}(t)} \left[\frac{\partial h_{\hat{u}}}{\partial t}(t,x) + \Omega(h_{\hat{u}}(t,\cdot))(x) \right] \cdot f(u) du = 0$$

by harmonicity of the functions h_v , using Proposition 3.1 with $V = [0, t_v]$.

To make this rigorous, all is needed is to prove Equation (25), by fixing some arbitrary $x \in \mathbb{N}^{*n}$ and $t_0 \ge 0$ and studying the ratio $[g(t_0 + \delta, x) - g(t_0, x)]/\delta$ as δ tends to zero. The monotonicity of the family of sets C(t) forces to distinguish the two cases $\delta > 0$ and (for $t_0 > 0$) $\delta < 0$. For the sake of shortness we will only present here the case $\delta > 0$, the other side being similar: both cases make a repeated use of the mean-value theorem and of Lebesgue's dominated convergence theorem.

To simplify notations, define for $u \in \mathcal{C}(t), t \ge 0$ and $x \in \mathbb{N}^{*n}$,

$$h(u,t,x) = h_{\hat{u}}(t,x) = e^{\varphi(\hat{u},t)} \prod_{i=1}^{n} (1 + \phi_i(\hat{u},t))^{x_i}.$$

Note that h inherits the derivability properties of φ with respect to t (the factor involving ϕ being \mathcal{C}^1 on \mathbb{R}): for fixed $u \in \mathcal{C}(0)$ and $x \in \mathbb{N}^{*n}$, h is \mathcal{C}^1 on $]-\infty, t_{\hat{u}}[$.

Let $x \in \mathbb{N}^{*n}$ and $t_0 \ge 0$ be fixed. For any positive δ , using the inclusion $\mathcal{D}(t_0 + \delta) \subset \mathcal{D}(t_0)$ one can write:

$$\frac{g(t_0+\delta,x)-g(t_0,x)}{\delta} = \int_{\mathcal{C}(t_0+\delta)} \frac{h(u,t_0+\delta,x)-h(u,t_0,x)}{\delta} f(u) du - \int_{\mathcal{C}(t_0)\backslash\mathcal{C}(t_0+\delta)} \frac{h(u,t_0,x)}{\delta} f(u) du.$$

Let first show that the first term tends to $\int_{\mathcal{C}(t_0)} \frac{\partial h}{\partial t}(u, t_0, x) f(u) du$ as δ tends to zero. Using the mean value theorem, since $h(u, \cdot, x)$ is \mathcal{C}^1 on $[0, t_0 + \delta]$ for $u \in \mathcal{C}(t_0 + \delta)$, this first term can be rewritten as

$$\int_{\mathcal{C}(t_0+\delta)} \frac{\partial h}{\partial t}(u, t_0 + p(u)\delta, x)f(u)du$$

for some $p(u) \in]0,1[$ depending on u, t_0, x and δ .

As δ goes to zero, $\frac{\partial h}{\partial t}(u, t_0 + p(u)\delta, x)f(u)$ tends to $\frac{\partial h}{\partial t}(u, t_0, x)f(u)$ and the indicator function of $C(t_0 + \delta)$ tends to the indicator function of $C(t_0)$ due to relation (20) which obviously extends to the sets C(t). The convergence of the first term will then result from Lebesgue's theorem by computing (we omit the variable (\hat{u}, t) under ϕ):

$$\frac{\partial h}{\partial t}(u,t,x) = e^{\varphi(\hat{u},t)} \sum_{i=1}^{n} (1+\phi_i)^{x_i-1} \prod_{j\neq i} (1+\phi_j)^{x_j} \left(\mu_i - \lambda_i \phi_i (1+\phi_i) + x_i \frac{\partial \phi_i}{\partial t}\right),$$

and then using the following domination. For $0 < \delta < 1$ and $u \in \mathbb{R}^{n-1}$,

$$\left|\frac{\partial h}{\partial t}(u,t_0+p(u)\delta,x)f(u)\mathbb{1}_{\mathcal{C}(t_0+\delta)}\right| \le k(t_0,x)\cdot|f(u)|\cdot\mathbb{1}_{\mathcal{C}(t_0)},$$

where the right hand side is integrable on \mathbb{R}^{n-1} , and $k(t_0, x)$ holds for:

$$\sup_{\mathcal{A}} e^{\varphi} \times \sup_{\mathcal{D}(t_0) \times [0, t_0+1]} \left| \sum_{i=1}^n (1+\phi_i)^{x_i-1} \prod_{j \neq i} (1+\phi_j)^{x_j} \left(\mu_i - \lambda_i \phi_i (1+\phi_i) + x_i \frac{\partial \phi_i}{\partial t} \right) \right|.$$

The convergence of the first term is thus proved.

We now prove that the second term vanishes as δ tends to 0. For any $u \in C(t_0) \setminus C(t_0 + \delta)$ there exists some index i (depending on u) such that $1 + \phi_i(\hat{u}, t_0) > 0$ while $1 + \phi_i(\hat{u}, t_0 + \delta) \leq 0$, and this implies by the mean value theorem that $0 < 1 + \phi_i(\hat{u}, t_0) \leq -\delta \frac{\partial \phi_i}{\partial t}(\hat{u}, t_0 + q(u)\delta)$ for some $q(u) \in]0, 1[$ depending on u, t_0 and δ . One can deduce the following upper bound, again assuming $0 < \delta < 1$:

$$\left| \int_{\mathcal{C}(t_0) \setminus \mathcal{C}(t_0+\delta)} \frac{h(u, t_0, x)}{\delta} f(u) du \right| \le k \int_{\mathcal{C}(t_0) \setminus \mathcal{C}(t_0+\delta)} |f(u)| du$$

where k is the following constant:

$$\sup_{\mathcal{A}} e^{\varphi} \times \max_{1 \le i \le n} \left\{ \sup_{\mathcal{D}(t_0) \times [0, t_0 + 1]} \left| \frac{\partial \phi_i}{\partial t} \right| \times \sup_{\mathcal{D}(t_0) \times \{t_0\}} \left| (1 + \phi_i)^{x_i - 1} \prod_{j \ne i} (1 + \phi_j)^{x_j} \right| \right\}.$$

The right hand side of the previous inequality converges to zero, again by Lebesgue's theorem, because f is integrable on the bounded set $C(t_0)$, and the sets $C(t_0) \setminus C(t_0 + \delta)$ decrease to \emptyset as δ decreases to zero, due to relation (20).

A.2. Change of Variables. The last step is now a change of variables in the harmonic function given by the integral (24), for a suitable choice of f so as to separate the time and space variables.

It informally consists in choosing as new variables the quantities $\pi_i(1+\phi_i(v,t))$ $(1 \leq i \leq n)$, changing the domain $\mathcal{D}(t)$ into $\mathcal{P} = \{v \in \mathbb{R}^n : v > 0 \text{ and } \sum_i^n v_i = 1\}$. Formally, it will be slightly more complicated due to integration with respect to Lebesgue's measure on subdomains of \mathbb{R}^{n-1} (the $\mathcal{C}(t)$'s), which forces to a round trip from \mathbb{R}^{n-1} through \mathbb{R}^n . So, to be correct, this change of variables will rather transform the domain $\mathcal{C}(t)$ into the following one:

(26)
$$\mathcal{S} = \left\{ u \in \mathbb{R}^{n-1} : u > 0 \text{ and } \sum_{i=1}^{n-1} u_i < 1 \right\}.$$

We need to introduce some additional notations. Denote by Δ the diagonal $n \times n$ square matrix having π_1, \ldots, π_n as its diagonal elements, by J the projection

$$J: \begin{array}{ccc} \mathbb{R}^n & \longrightarrow & \mathbb{R}^{n-1} \\ (v_1, \dots, v_n) & \longmapsto & (v_1, \dots, v_{n-1}) \end{array}$$

and by \mathcal{K} , the hyperplane of \mathbb{R}^n defined by

$$\mathcal{K} = \left\{ v \in \mathbb{R}^n : \sum_{i=1}^n v_i = 1 \right\}.$$

The hyperplane \mathcal{K} corresponds to \mathbb{R}^{n-1} through the one to one affine transformation (analogous to H from \mathbb{R}^{n-1} to \mathcal{H}):

$$K: \begin{array}{ccc} \mathbb{R}^{n-1} & \longrightarrow & \mathcal{K} \\ u & \longmapsto & \tilde{u} = (u, 1 - \sum_{i=1}^{n-1} u_i) \end{array}$$

Notice that the inverse mapping of K (resp. H) is given by the restriction of J to \mathcal{K} (resp. \mathcal{H}). The announced change of variables is given by the *t*-depending transformation

$$\Psi_t: \begin{array}{ccc} \mathbb{R}^{n-1} & \longrightarrow & \mathbb{R}^{n-1} \\ u & \longmapsto & J\Delta(P_{-t}Hu + \mathbb{1}) \end{array}$$

The next lemma shows that Ψ_t can be considered for a change of variables:

LEMMA A.3. For any $t \geq 0$, Ψ_t is a one to one affine transformation on \mathbb{R}^{n-1} which inverse mapping is given by

$$\Psi_t^{-1}(u) = JP_t(\Delta^{-1}Ku - \mathbb{1})$$

and which Jacobian is $Jac(\Psi_t) = e^{\theta t} \prod_{i=1}^{n-1} \pi_i$. Moreover, $\Psi_t(\mathcal{C}(t)) = \mathcal{S}$.

PROOF. Since Ψ_t is clearly an affine transformation in \mathbb{R}^{n-1} , its Jacobian is the one of its linear part $J\Delta P_{-t}H$. Now $J\Delta = \Delta' J$, where Δ' is the diagonal $(n-1) \times (n-1)$ square matrix having π_1, \ldots, π_{n-1} as its diagonal elements, so that

$$\operatorname{Jac}(J\Delta P_{-t}H) = \left(\prod_{i=1}^{n-1} \pi_i\right) \operatorname{Jac}(JP_{-t}H) = \left(\prod_{i=1}^{n-1} \pi_i\right) \operatorname{Jac}(P_{-t}) = e^{\theta t} \prod_{i=1}^{n-1} \pi_i.$$

The second equality results from the facts that J restricted to \mathcal{H} coincides with H^{-1} and that \mathcal{H} is generated by the first n-1 eigenvectors of P_{-t} , so that \mathcal{H} is invariant under P_{-t} , which restriction to \mathcal{H} has Jacobian $\prod_{i=1}^{n-1} e^{-\theta_i t} = \prod_{i=1}^n e^{-\theta_i t} =$ $\operatorname{Jac}(P_{-t})(=e^{\theta t})$ since $\theta_n = 0$. In particular $\operatorname{Jac}(\Psi_t) \neq 0$, so that Ψ_t is invertible. The formula for Ψ^{-1} easily results from the fact that $\Delta(\cdot + 1)$ maps \mathcal{H} onto \mathcal{K} and that the inverse mapping of K is given by the restriction of J to \mathcal{K} .

Now using $C(t) = \{u \in \mathbb{R}^{n-1} : \mathbb{1} + P_{-t}Hu > 0\}$ together with the facts that Δ preserves the relation v > 0 and (again) that $\Delta(\cdot + \mathbb{1})$ maps \mathcal{H} onto \mathcal{K} , one gets that $\Psi_t(\mathcal{C}(t))$ is included in $J(\{v \in \mathcal{K} : v > 0\}) = \mathcal{S}$. Equality results from a similar argument for $\Psi_t^{-1}(\mathcal{S}) \subset C(t)$.

The transformation Ψ_t hence corresponds to the two following diagrams:

$$\begin{array}{ccc} \mathcal{H} \subset \mathbb{R}^n \xrightarrow{\Delta(P_{-t} \cdot + 1)} \mathcal{K} \subset \mathbb{R}^n & \mathcal{H} \subset \mathbb{R}^n \xrightarrow{P_t(\Delta^{-1} \cdot - 1)} \mathcal{K} \subset \mathbb{R}^n \\ & & \\ H & & \\ & & \\ \mathcal{L}(t) \subset \mathbb{R}^{n-1} \xrightarrow{\Psi_t} \mathcal{S} \subset \mathbb{R}^{n-1} & & \\ \mathcal{L}(t) \subset \mathbb{R}^{n-1} \xleftarrow{\Psi_t} \mathcal{S} \subset \mathbb{R}^{n-1} \end{array}$$

All is left now is to choose for (24) a family of locally integrable functions in \mathbb{R}^{n-1} which behave nicely with respect to the change of variables Ψ_t . It will be given by the functions $f^{\alpha-1}$ for positive α 's, where, for $u \in \mathbb{R}^{n-1}$ and $v \in \mathbb{R}^n$,

$$f(u) = \psi(Hu) = \psi(\hat{u})$$
 and $\psi(v) = \prod_{i=1}^{n-1} |(\omega^{-1}v)_i|$

Here, for any $v \in \mathbb{R}^n$ and $1 \leq i \leq n$, v_i denotes the *i*th coordinate of v, so that the $(\omega^{-1}v)_i$'s $(1 \leq i \leq n-1)$ are the first n-1 coordinates of v in the base $(\omega_1, \ldots, \omega_n)$ of eigenvectors of Q. As will become clear in (29), the next lemma establishes the property of ψ that makes it behave nicely with respect to the change of variables given by Ψ_t , by isolating the dependency in time in a separate factor:

LEMMA A.4. For any $t \ge 0$ and $v \in \mathbb{R}^n$, $\psi(P_t v) = e^{-\theta t} \psi(v)$.

PROOF. This result stems from diagonalizing P_t as $\omega^{-1}P_t\omega = e^{-t\Theta}$, where Θ is the diagonal $n \times n$ square matrix having $\theta_1, \ldots, \theta_n$ as its diagonal elements. This readily gives $\psi(P_t v) = \prod_{i=1}^{n-1} |(e^{-t\Theta}\omega^{-1}v)_i| = e^{-\theta t}\psi(v)$.

The main technical point is to establish that $f^{\alpha-1}$ is locally integrable: the next lemma provides in addition a useful upper bound.

LEMMA A.5. The function $f^{\alpha-1}$ is locally integrable on \mathbb{R}^{n-1} for any $\alpha > 0$. Moreover, for any compact set $T \subset \mathbb{R}^{n-1}$,

(27)
$$\sup_{0<\alpha\leq 1} \left(\alpha^n \int_T f(u)^{\alpha-1} du\right) < +\infty.$$

PROOF. If $\alpha \geq 1$, $f^{\alpha-1}$ is continuous on \mathbb{R}^{n-1} , hence locally integrable. So consider only the case when $0 < \alpha < 1$

If the matrix ω has real coefficients (it can be chosen such when the eigenvalues θ_j of Q are real, which is in particular the case for a reversible Q), $f^{\alpha-1}$ is easily shown to be integrable on any compact set T of \mathbb{R}^{n-1} by operating the change of variable:

$$\begin{array}{cccc} \mathbb{R}^{n-1} & \longrightarrow & \mathbb{R}^{n-1} \\ u & \longmapsto & \left((\omega^{-1}Hu)_i \right)_{1 < i < n-1} = J\omega^{-1}Hu \end{array}$$

which is linear and one to one, and transforms $\int_T f(u)^{\alpha-1} du$ into the integral over some compact subset of \mathbb{R}^{n-1} of the locally integrable function $\prod_{i=1}^{n-1} |u_i|^{\alpha-1}$ (up to the Jacobian constant factor). Then by considering A large enough so that $J\omega^{-1}H(T) \subset [-A, A]^{n-1}$ and $A \ge 1$, (27) is obtained from the fact that

$$\int_{[-A,A]^{n-1}} \prod_{i=1}^{n-1} |u_i|^{\alpha-1} \, du = (2A^{\alpha})^{n-1} \alpha^{-(n-1)} \le (2A)^{n-1} \alpha^{-n}$$

This is not directly possible when ω has non-real coefficients. In this case we can show that $f^{\alpha-1}$ is upper bounded by $\prod_{i=1}^{n-1} |(Lu)_i|^{\alpha-1}$ for some invertible $(n-1) \times (n-1)$ square matrix L with real coefficients. The change of variables $u \mapsto Lu$ is then possible, showing (27) in this case similarly as before, which implies the local integrability of $f^{\alpha-1}$ for $0 < \alpha < 1$.

Since $\alpha < 1$ this amounts to lower bounding f by $\prod_{i=1}^{n-1} |(Lu)_i|$.

Call C the complex invertible $(n-1) \times (n-1)$ square matrix associated to the linear mapping $J\omega^{-1}H$ on \mathbb{R}^{n-1} , so that $f(u) = \prod_{j=1}^{n-1} |(Cu)_j|$, and write C = A + iB where A and B are real square matrices. For any $p \in [0,1]$, $u \in \mathbb{R}^{n-1}$ and $j \in \{1, \ldots, n-1\}$, the following inequalities hold:

$$|(Cu)_j| \ge \max\{|(Au)_j|, |(Bu)_j|\} \ge |p(Au)_j + (1-p)(Bu)_j| = \left|\left((pA + (1-p)B)u\right)_j\right|.$$

All is left now is to prove the existence of some $p \in [0, 1]$ such that pA + (1-p)B is invertible. It is done through considering the degree n-1 polynomial with complex variable: det(A+zB), which is non zero at z = i (since det $C \neq 0$), hence not equal to the null polynomial. It then cannot be zero on the whole real interval [0, 1], which gives the result.

PROOF OF THEOREM 3.1. The previous lemma shows that $f^{\alpha-1}$ is a suitable function to plug in (24): since $\varphi(v,t) = \varphi_0(P_{-t}v)$ for $(v,t) \in \mathcal{A}$, rewriting (24) and using the definition of f gives

$$g(t,x) = \int_{\mathcal{C}(t)} e^{\varphi(\hat{u},t)} \prod_{i=1}^{n} (1+\phi_i(\hat{u},t))^{x_i} f(u)^{\alpha-1} du$$
$$= \int_{\mathcal{C}(t)} e^{\varphi_0(P_{-t}Hu)} \prod_{i=1}^{n} (1+P_{-t}Hu)^{x_i}_i \psi(P_t P_{-t}Hu)^{\alpha-1} du.$$

Expressing $P_{-t}Hu$ through $\Psi_t u$ for $u \in \mathbb{R}^{n-1}$, one gets, since K is the inverse of J restricted to \mathcal{K} and $\Delta(P_{-t}Hu + 1) \in \mathcal{K}$,

(28)
$$P_{-t}Hu = \Delta^{-1}K\Psi_t u - \mathbb{1}, \quad u \in \mathcal{C}(t)$$

so that operating the change of variables given by Ψ_t yields by Lemma A.3

(29)
$$g(t,x) = e^{\theta t} \prod_{i=1}^{n-1} \pi_i \int_{\mathcal{S}} e^{\varphi_0(\Delta^{-1}\tilde{u}-1)} \prod_{i=1}^n \left(\Delta^{-1}\tilde{u}\right)_i^{x_i} \psi \left(P_t(\Delta^{-1}\tilde{u}-1)\right)^{\alpha-1} du.$$

Since $\omega_n = 1$, we have $\psi(v - 1) = \psi(v)$ for any $v \in \mathbb{R}^n$, hence $\psi\left(P_t(\Delta^{-1}\tilde{u} - 1)\right) = \psi\left(P_t\Delta^{-1}\tilde{u}\right) = e^{-\theta t}\psi(\Delta^{-1}\tilde{u})$, where the last equality comes from Lemma A.4. The following function g', which only differs from g by a multiplicative factor, is thus

again space-time harmonic:

$$\begin{split} g'(t,x) &= e^{\theta t} \int_{\mathcal{S}} G(\tilde{u}) \prod_{i=1}^{n} \left(\Delta^{-1} \tilde{u} \right)_{i}^{x_{i}} \left(e^{-\theta t} \psi(\Delta^{-1} \tilde{u}) \right)^{\alpha - 1} du \\ &= e^{-\alpha \theta t} \int_{\mathcal{S}} G(\tilde{u}) \prod_{i=1}^{n} \left(\frac{\tilde{u}_{i}}{\pi_{i}} \right)^{x_{i}} \psi(\Delta^{-1} \tilde{u})^{\alpha - 1} du, \end{split}$$

where we have defined

$$G(v) = e^{\varphi_0(\Delta^{-1}v - \mathbb{1})}, \quad v \in \mathcal{P}.$$

Hence defining F as

$$F(v) = \psi(\Delta^{-1}v), \quad v \in \mathbb{R}^n$$

yields exactly the local martingale of Theorem 3.1. The second expression (2) is easily obtained. All one needs to do to complete the proof of Theorem 3.1 is to check the announced properties of these two functions F and G.

First, G is continuous and bounded on \mathcal{P} since, by Lemma A.2, φ_0 is continuous and bounded above on $\mathcal{D}(0)$ (if $v \in \mathcal{P}$, then $\Delta^{-1}v - \mathbb{1} \in \mathcal{H}$ and $\Delta^{-1}v > 0$, so that $\Delta^{-1}v - \mathbb{1} \in \mathcal{D}(0)$).

Moreover, F is clearly positive and continuous on \mathbb{R}^n , and thus bounded on the bounded subset \mathcal{P} of \mathbb{R}^n , and so (3) is the only property left to be checked.

Relation (28) and the fact that $\psi(\cdot - 1) = \psi(\cdot)$, together with the definitions of F and f, yield that $F(\tilde{u}) = f(\Psi_0^{-1}u)$ according to the following steps:

 $F(\tilde{u}) = \psi(\Delta^{-1}Ku) = \psi(\Delta^{-1}K\Psi_0(\Psi_0^{-1}u) - 1) = \psi(H\Psi_0^{-1}u) = f(\Psi_0^{-1}u).$

It follows by the change of variables induced by Ψ_0 that, for $\alpha \leq 1$,

$$\int_{\mathcal{S}} F(\tilde{u})^{\alpha-1} du = \int_{\mathcal{S}} f(\Psi_0^{-1}u)^{\alpha-1} du = \operatorname{Jac}(\Psi_0) \int_T f(u)^{\alpha-1} du,$$

where $T = \overline{\Psi_0^{-1}(S)}$. The property (3) then follows using Lemma A.5.

References of Chapter I

- [AK98] Krishna B. Athreya and Hye-Jeong Kang. Some limit theorems for positive recurrent branching Markov chains: I. Advances in Applied Probability, 30(3):711–722, September 1998.
- [Bor05] Sem Borst. User level performance of channel-aware scheduling algorithms in wireless data networks. *IEEE/ACM Transactions on Network*ing, 13(3):636–647, June 2005.
- [BPH06] Sem Borst, Alexandre Proutière, and Nidhi Hedge. Capacity of wireless data networks with intra- and inter-cell mobility. In *Proceedings of IEEE* INFOCOM 2006, pages 1058–1069, Barcelona, Spain, April 2006.
- [Bra96a] Maury Bramson. Convergence to equilibria for fluid models of FIFO queueing networks. Queueing Systems: Theory and Applications, 22(1– 2):5–45, 1996.
- [Bra96b] Maury Bramson. Convergence to equilibria for fluid models of head-of-theline proportional processor sharing queueing networks. *Queueing Systems: Theory and Applications*, 23(1–4):1–26, 1996.
- [CP07] Pietro Caputo and Gustavo Posta. Entropy dissipation estimates in a zero-range dynamics. Probability Theory and Related Fields, 139:65–87, September 2007.
- [CY01] Hong Chen and David D. Yao. Fundamentals of queueing networks. Springer-Verlag, New York, 2001. Performance, asymptotics, and optimization, Stochastic Modelling and Applied Probability.
- [Dai95] J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. Annals of Applied Probability, 5(1):49–77, February 1995.
- [DW94] P. Dupuis and R. J. Williams. Lyapounov functions for semimartingale reflecting Brownian motions. Annals of Applied Probability, 22(2):680– 702, 1994.
- [FRT99] Christine Fricker, Philippe Robert, and Danielle Tibi. On the rates of convergence of Erlang's model. *Journal of Applied Probability*, 36(4):1167– 1184, 1999.
- [GT01] M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In *Proceedings of IEEE INFOCOM 2001*, volume 3, pages 1360–1369, 2001.
- [Rob03] Philippe Robert. Stochastic Networks and Queues. Stochastic Modelling and Applied Probability Series. Springer-Verlag, New York, 2003. xvii+398 pp.
- [RS92] A. N. Rybko and A. L. Stolyar. On the ergodicity of random processes that describe the functioning of open queueing networks. *Problems on Information Transmission*, 28(3):3–26, 1992.

CHAPTER II

Stability Properties of Linear File-Sharing Networks

Contents	\mathbf{C}	ont	en	\mathbf{ts}
----------	--------------	-----	----	---------------

1. Introduction	83
2. Analysis of the Single-Chunk Network	85
3. Yule Processes with Deletions	91
4. Analysis of the Multi-Chunk Network	97
Appendix A. Proof of Proposition 3.3	104
References of Chapter II	107

1. Introduction

File-sharing networks are distributed systems used to disseminate information among a subset of the nodes of the Internet (overlay network). The general simple principle is the following: once a node of the system has retrieved a file it becomes a server for this file. The advantage of this scheme is that it disseminates information in a very efficient way as long as the number of servers is growing rapidly. The growth of the number of servers is not necessarily without bounds since a node having this file may stop being a server after some time. These schemes have been used for some time now in peer-to-peer systems such as BitTorrent or Emule, for example to distribute large files over the Internet.

An improved version of this principle consists in splitting the original file into several pieces (called "chunks") so that a given node can retrieve simultaneously several chunks of the same file from different servers. In this case, the rate to get a given file may thus increase significantly. At the same time, the global capacity of the file-sharing system is also increased since a node becomes a server of a chunk as soon as it has retrieved it and not only when it has the whole file. This improvement has interesting algorithmic implications since each node has to establish a matching between chunks and servers. Strategies to maximize the global efficiency of the file sharing systems have to be devised. See for instance Massoulié and Vojnović [**MV05**], Bonald *et al.* [**BMM**⁺08] and Massoulié and Twigg [**MT08**].

The efficiency of these systems can be considered from different points of view.

Transient behavior: A new file is owned by one node, given there are potentially N other nodes interested by it, how long does it take so that a given node retrieves it ? significant fraction $\alpha \in (0, 1]$ of the N nodes

retrieve it ? See Yang and de Veciana [YdV04], and Chapters III and IV (corresponding respectively to [SRG08] and [RS09]).

Stationary behavior: A constant flow of requests enters, is the capacity of the file-sharing system sufficient to cope with this flow ?

In this chapter, the stationary behavior is investigated in a stochastic context: arrival times are random as well as chunk transmission times. In this setting mathematical studies are quite scarce, see Qiu and Srikant [QS04], Simatos *et al.* [SRG08], Susitaival *et al.* [SAV06] and references therein. A simple strategy to disseminate chunks is considered: chunks are retrieved sequentially and a given node can be the server of only the last chunk it got. See Massoulié and Vojnović [MV05] and Parvez *et al.* [PWMC08] for a detailed motivation of this situation.

In this chapter, the sequential scheme for disseminating a file that is divided into n chunks is analyzed. New requests arrive according to a Poisson process at rate λ , and become downloaders of chunk 1. Users who have obtained chunks $1, \ldots, k$ act simultaneously as uploaders of chunk k and downloaders of chunk k+1, and the users who have all the chunks leave the network at rate ν . The transmission rate of chunk k is denoted by μ_k , and x_k is the number of users having obtained chunks $1, \ldots, k$. In this way, the total transmission rate of chunk k in the network is $\mu_k x_k$. The flow of users can be modeled as the linear network depicted in Figure 1.



FIGURE 1. Transition rates of the linear network outside boundaries.

The main problem analyzed in the chapter is the determination of a constant λ^* such that if $\lambda < \lambda^*$ [resp. $\lambda > \lambda^*$], then the associated Markov process is ergodic [resp. transient]. As it will be seen, the constant λ^* may be infinite in some cases so that the file-sharing network is always stable independently of the value of λ . The main technical difficulty to prove stability/instability results for this class of stochastic networks is that, except for the input, the Markov process has unbounded jump rates, in fact proportional to one of the coordinates of the current state. Note that loss networks have also this characteristic but in this case, the stability problem is trivial since the state space is finite. See Kelly [Kel91].

Fluid Limits for File-Sharing Networks. Classically, to analyze the stability properties of stochastic networks, one can use the limits of a scaling of the Markov process, the so-called fluid limits. The scaling consists in speeding up time by the norm ||x|| of the initial state x, by scaling the state vector by 1/||x|| and by letting ||x|| go to infinity. See Bramson [Bra08], Chen and Yao [CY01] and Robert [Rob03] for example. This scaling is, however, better suited to "locally additive" processes, that is, Markov processes that behave locally as random walks. Since the transition rates are unbounded, it may occur that the corresponding fluid limits have discontinuities; this complicates a lot the analysis of a possible limiting dynamical system. Roughly speaking, this is due to the fact that, because of the unbounded transition rates, events occur on the time scale $t \mapsto t \log ||x||$ instead of

 $t \mapsto ||x||t$. See the case of the $M/M/\infty$ queue in Chapter 9 of Robert [**Rob03**], and Simatos and Tibi [**ST09**] for a discussion of this phenomenon in a related context.

A "fluid scaling" is nevertheless available for file-sharing networks. A possible description for a possible candidate $(x_i(t))$ for this limiting picture would satisfy the following differential equations,

(1)
$$\begin{cases} \dot{x}_0(t) = \lambda - \mu_1 x_1(t), \\ \dot{x}_i(t) = \mu_i x_i(t) - \mu_{i+1} x_{i+1}(t), \quad 1 \le i \le n-1, \\ \dot{x}_n(t) = \mu_n x_n(t) - \nu x_n(t). \end{cases}$$

For the sake of simplicity the behavior at the boundaries $\{x : x_i = 0\}, i \geq 1$ is ignored in the above equations. This has been, up to now, one of the main tools to investigate mathematical models of file-sharing networks. See Qiu and Srikant [QS04], Núñez-Queija and Prabhu [NQP08] for example. In the context of loss networks, an analogous limiting picture can be rigorously justified when the input rates and buffer sizes are scaled by some N and the state variable by 1/N. This scaling is not useful here, since the problem is precisely of determining the values of λ for which the associated Markov is ergodic whereas in the above scaling λ is scaled. From this point of view Equations (1) are therefore quite informal. They can nevertheless give some insight into the qualitative behavior of these networks but they cannot apparently be used to prove stability results. Their interpretation near boundaries is in particular not clear.

Interacting Branching Processes. Since scaling techniques do not apply here, one needs to resort to different techniques to study stability: coupling the linear file-sharing network with interacting branching processes is a key idea. For $i \ge 1$, without the departures the process $(X_i(t))$ would be a branching process where individuals give birth to one child at rate μ_i . This description of such a file-sharing system as a branching process is quite natural. It has been used to analyze the transient behavior of these systems. See Yang and de Veciana [YdV04], Dang *et al.* [DPM07] and Simatos *et al.* [SRG08]. A departure for $(X_i(t))$ can be seen as a death of an individual of class *i* and at the same time as a birth of an individual of class *i* and at the same time as a system of interacting branching processes with a constant input rate λ .

To tackle the general problem of stability, several key ingredients are used in this chapter: Lyapunov functions, coupling arguments and precise estimations of the growth of a branching process killed by another branching process. As it will be seen, several results used come from the branching process formulation of the stochastic model. In particular Section 3 is devoted to the derivation of results concerning killed branching processes. The stability properties of networks with a single-chunk file are analyzed in detail in Section 2. In Section 4, file-sharing networks with n chunks are studied and the case n = 2 is investigated thoroughly.

2. Analysis of the Single-Chunk Network

This section is devoted to the study of a class of two-dimensional Markov jump processes $(X_0(t), X_1(t))$, for $x = (x_0, x_1) \in \mathbb{N}^2$, the corresponding Q-matrix Ω_r is given by

(2)
$$\begin{cases} \Omega_r[(x_0, x_1), (x_0 + 1, x_1)] &= \lambda, \\ \Omega_r[(x_0, x_1), (x_0 - 1, x_1 + 1)] &= \mu r(x_0, x_1)(x_1 \vee 1) \mathbb{1}_{\{x_0 > 0\}}, \\ \Omega_r[(x_0, x_1), (x_0, x_1 - 1)] &= \nu x_1, \end{cases}$$

where $x \mapsto r(x)$, referred to as the *rate function*, is some fixed function on \mathbb{N}^2 with values in [0, 1] and $n \lor m$ denotes $\max(n, m)$ for $n, m \in \mathbb{N}^2$. This corresponds to a more general model than the linear file-sharing network of Figure 1 in the case n = 1, where for the sake of simplicity μ_1 is noted μ in this section.

From a modeling perspective, this Markov process describes the following system. Requests for a single file arrive with rate λ , the first component $X_0(t)$ is the number of requests which did not get the file, whereas the second component is the number of requests having the file and acting as servers until they leave the file-sharing network. The constant μ can be viewed as the file transmission rate, and ν as the rate at which servers having all chunks leave. The term $r(x_0, x_1)$ describes the interaction of downloaders and uploaders in the system. The term $x_1 \vee 1$ can be interpreted so that there is one server permanent server in the network, which is contacted if there are no other uploader nodes in the system. A related system where there is always one permanent server for the file can be modeled by replacing the term $x_1 \vee 1$ by $x_1 + 1$. See the remark at the end of this section.

Several related examples of this class of models have been recently investigated. The case

$$(x_0, x_1) = \frac{x_0}{x_0 + x_1}$$

is considered in Núñez-Queija and Prabhu [**NQP08**] and and also in the paper by Massoulié Vojnović [**MV05**]; in this case the downloading time of the file is neglected. Susitaival *et al.* [**SAV06**] analyzes the rate function r(x)

$$r(x_0, x_1) = 1 \land \left(\alpha \frac{x_0}{x_1}\right)$$

with $\alpha > 0$ and $a \wedge b$ denotes $\min(a, b)$ for $a, b \in \mathbb{R}$. This model allows to take into account that a request cannot be served by more than one server. See also Qiu and Srikant [QS04].

With a slight abuse of notation, for $0 < \delta \leq 1$, the matrix Ω_{δ} will refer to the case when the function r is identically equal to δ . Note that the boundary condition $x_1 \vee 1$ for departures from the first queue prevents the second coordinate from ending up in the absorbing state 0. Other possibilities are discussed at the end of this section. In the following $(X^r(t)) = (X_0^r(t), X_1^r(t))$ [resp. $(X^{\delta}(t))$] will denote a Markov process with Q-matrix Ω_r [resp. Ω_{δ}].

Free Process. For $\delta > 0$, Q_{δ} denotes the following Q-matrix

r

(3)
$$\begin{cases} Q_{\delta}[(y_0, y_1), (y_0 + 1, y_1)] &= \lambda, \\ Q_{\delta}[(y_0, y_1), (y_0 - 1, y_1 + 1)] &= \mu \delta(y_1 \vee 1), \\ Q_{\delta}[(y_0, y_1), (y_0, y_1 - 1)] &= \nu y_1. \end{cases}$$

The process $(Y^{\delta}(t)) = (Y_0^{\delta}(t), Y_1^{\delta}(t))$, referred to as the free process, will denote a Markov process with *Q*-matrix Q_{δ} . Note that the first coordinate Y_0^{δ} may become negative. The second coordinate $(Y_1^{\delta}(t))$ of the free process is a classical birth-and-death process. It is easily checked that if ρ_{δ} defined as $\delta \mu / \nu$ is such that $\rho_{\delta} < 1$,

then $(Y_1^\delta(t))$ is an ergodic Markov process converging in distribution to $Y_1^\delta(\infty)$ and that

(4)
$$\lambda^*(\delta) \stackrel{\text{def.}}{=} \nu \mathbb{E}(Y_1^{\delta}(\infty)) = \mu \mathbb{E}(Y_1^{\delta}(\infty) \vee 1) = \frac{\delta \mu}{(1 - \rho_{\delta})(1 - \log(1 - \rho_{\delta}))}$$

When $\rho_{\delta} > 1$, then the process $(Y^{\delta}(t))$ converges almost surely to infinity. In the sequel $\lambda^*(1)$ is simply denoted λ^* .

In the following it will be assumed, Condition (C) below, that the rate function r converges to 1 as the first coordinate goes to infinity; as will be seen, the special case $r \equiv 1$ then plays a special role, and so before analyzing the stability properties of $(X^r(t))$, one begins with an informal discussion when the rate function r is identically equal to 1. Since the departure rate from the system is proportional to the number of requests/servers in the second queue, a large number of servers in the second queue gives a high departure rate, irrespectively of the state of the first queue. The input rate of new requests being constant, the real bottleneck with respect to stability is therefore when the first queue is large. The interaction of the two processes $(X_0^1(t))$ and $(X_1^1(t))$ is expressed through the indicator function of the set $\{X_0^1(t) > 0\}$. The second queue $(X_1^1(t))$ locally behaves like the birth-and-death process $(Y_1^1(t))$ as long as $(X_0^1(t))$ is away from 0. The two cases $\rho_1 > 1$ and $\rho_1 < 1$ are considered.

If $\rho_1 > 1$, i.e., $\mu > \nu$, the process $(X_1^1(t))$ is a transient process as long as the first coordinate is non-zero. Consequently, departures from the second queue occur faster and faster. Since, on the other hand, arrivals occur at a steady rate, departures eventually outpace arrivals. The fact that the second queue grows when $(X_0(t))$ is away from 0 stabilizes the system independently of the value of λ , and so the system should be stable for any $\lambda > 0$.

If $\rho_1 < 1$, and as long as $(X_0(t))$ is away from 0, the coordinate $(X_1^1(t))$ locally behaves like the ergodic Markov process $(Y_1^1(t))$. Hence if $(X_0^1(t))$ is non-zero for long enough, the requests in the first queue see in average $\mathbb{E}(Y_1^1(\infty) \vee 1)$ servers which work at rate μ . Therefore, the stability condition for the first queue should be

$$\lambda < \mu \mathbb{E}(Y_1^1(\infty) \lor 1) = \lambda$$

where $\lambda^* = \lambda^*(1)$ is defined by Equation (4). Otherwise if $\lambda > \lambda^*$, the system should be unstable.

Markovian Notations. In the following, one will use the following convention, if (U(t)) is a Markov process, the index u of $\mathbb{P}_u((U(t)) \in \cdot)$ will refer to the initial condition of this Markov process.

Transience and Recurrence Criteria for $(X^r(t))$.

PROPOSITION 2.1 (Coupling). If $X^r(0) = Y^1(0) \in \mathbb{N}^2$, there exists a coupling of the processes $(X^r(t))$ and $(Y^1(t))$ such that the relation

(5)
$$X_0^r(t) \ge Y_0^1(t) \text{ and } X_1^r(t) \le Y_1^1(t),$$

holds for all $t \ge 0$ and for any sample path.

For any $0 \leq \delta \leq 1$, if

$$\tau_{\delta} = \inf\{t \ge 0 : r(X^{r}(t)) \le \delta\} \text{ and } \sigma = \inf\{t \ge 0 : X_{0}^{r}(t) = 0\},\$$

and if $X^1(0) = Y^{\delta}(0) \in \mathbb{N}^2$ then there exists a coupling of the processes $(X^r(t))$ and $(Y^{\delta}(t))$ such that, for any sample path, the relation

$$X_0^r(t) \le Y_0^{\delta}(t) \text{ and } X_1^r(t) \ge Y_1^{\delta}(t)$$

holds for all $t \leq \tau_{\delta} \wedge \sigma$.

(6)

PROOF. Let $X^r(0) = (x_0, x_1)$ and $Y^1(0) = (y_0, y_1)$ be such that $x_0 \ge y_0$ and $x_1 \le y_1$, one has to prove that the processes $(X^r(t))$ and $(Y^1(t))$ can be constructed such that Relation (5) holds at the time of the next jump of one of them. See Leskelä **[Les09]** for the existence of couplings using analytical, nonconstructive techniques.

The arrival rates in the first queue are the same for both processes. If $x_1 < y_1$, a departure from the second queue for $(Y^1(t))$ or $(X^r(t))$ preserves the order relation (5) and if $x_1 = y_1$, this departure occurs at the same rate for both processes and thus the corresponding instant can be chosen at the same (exponential) time. For the departures from the first to the second queue, the departure rate for $(X^r(t))$ is $\mu r(x_0, x_1)(x_1 \vee 1)\mathbb{1}_{\{x_0>0\}} \leq \mu(y_1 \vee 1)$ which is the departure rate for $(Y^1(t))$, hence the corresponding departure instants can be taken in the reverse order so that Relation (5) also holds at the next jump instant. The first part of the proposition is proved.

The rest of the proof is done in a similar way: The initial states $X^r(0) = (x_0, x_1)$ and $Y^{\delta}(0) = (y_0, y_1)$ are such that $x_0 \leq y_0$ and $x_1 \geq y_1$. With the killing of the processes at time $\tau_{\delta} \wedge \sigma$ one can assume additionally that $x_0 \neq 0$ and that the relation $r(x_0, x_1) \geq \delta$ holds; Under these assumptions one can check by inspecting the next transition that (6) holds. The proposition is proved.

PROPOSITION 2.2. Under the condition $\mu < \nu$, the relation

$$\liminf_{t \to +\infty} \frac{X_0^r(t)}{t} \ge \lambda - \lambda^*$$

holds almost surely. In particular, if $\mu < \nu$ and $\lambda > \lambda^*$, then the process $(X^r(t))$ is transient.

PROOF. By Proposition 2.1, one can assume that there exists a version of $(Y^1(t))$ such that $X_0^r(0) = Y_0^1(0)$ and the relation $X_0^r(t) \ge Y_0^1(t)$ holds for any $t \ge 0$. From Definition (3) of the *Q*-matrix of $(Y^1(t))$, one has, for $t \ge 0$,

$$Y^{1}(t) = Y^{1}(0) + \mathcal{N}_{\lambda}(t) - A(t),$$

where $(\mathcal{N}_{\lambda}(t))$ is a Poisson process with parameter λ and (A(t)) is the number of arrivals (jumps of size 1) for the second coordinate $(Y_1^1(t))$: in particular

$$\mathbb{E}(A(t)) = \mu \mathbb{E}\left(\int_0^t Y_1^1(s) \vee 1 \, ds\right).$$

Since $(Y_1^1(t))$ is an ergodic Markov process under the condition $\mu < \nu$, the ergodic theorem in this setting gives that

$$\lim_{t \to +\infty} \frac{1}{t} A(t) = \lim_{t \to +\infty} \frac{1}{t} \mathbb{E}(A(t)) = \mu \mathbb{E}\left(Y_1^1(\infty) \lor 1\right) = \lambda^*,$$

by Equation (4), hence $(Y_0^1(t)/t)$ converges almost surely to $\lambda - \lambda^*$. The proposition is proved.

The next result establishes the ergodicity result of this section.

PROPOSITION 2.3. If the rate function r is such that, for any $x_1 \in \mathbb{N}$,

(C)
$$\lim_{x_0 \to +\infty} r(x_0, x_1) = 1,$$

and if $\mu \geq \nu$, or if $\mu < \nu$ and $\lambda < \lambda^*$ with

(7)
$$\lambda^* = \frac{\mu}{(1-\rho)(1-\log(1-\rho))},$$

and $\rho = \mu/\nu$, then $(X^r(t))$ is an ergodic Markov process.

Note that Condition (C) is satisfied for the functions r considered in the models considered by Núñez-Queija and Prabhu [NQP08] and in Susitaival *et al.* [SAV06]. See above.

PROOF. If $x = (x_0, x_1) \in \mathbb{R}^2$, |x| denotes the norm of x, $|x| = |x_0| + |x_1|$. The proof uses Foster's criterion as stated in Robert [**Rob03**, Theorem 9.7]. If there exist constants K_0 , K_1 , t_0 , t_1 and $\eta > 0$ such that, for $x = (x_0, x_1) \in \mathbb{N}^2$,

(8)
$$\mathbb{E}_{(x_0,x_1)}(|X^r(t_1)| - |x|) \le -t_1, \text{ if } x_1 \ge K_1$$

(9)
$$\mathbb{E}_{(x_0, x_1)}(|X^r(t_0)| - |x|) \le -\eta t_0$$
, if $x_0 \ge K_0$ and $x_1 < K_1$,

then the Markov process $(X^r(t))$ is ergodic.

Relation (8) is straightforward to establish: if $x_1 \ge K_1$, one gets, by considering only K_1 of the x_1 initial servers in the second queue and the Poisson arrivals, that

$$\mathbb{E}_{(x_0,x_1)}(|X^r(1)| - |x|) \le \lambda - K_1(1 - e^{-\nu}),$$

hence it is enough to take $t_1 = 1$ and $K_1 = (\lambda + 1)/(1 - e^{-\nu})$ to have Relation (8).

One has therefore to establish Inequality (9). Let τ_{δ} and σ be the stopping times introduced in Proposition 2.1, one first proves an intermediate result: for any t > 0 and any $x_1 \in \mathbb{N}$,

(10)
$$\lim_{x_0 \to +\infty} \mathbb{P}_{(x_0, x_1)}(\sigma \wedge \tau_{\delta} \le t) = 0.$$

Fix $x_1 \in \mathbb{N}$ and $t \ge 0$: for $\varepsilon > 0$, there exists D_1 such that

$$\mathbb{P}_{x_1}\left(\sup_{0\le s\le t}Y_1^1(s)\ge D_1\right)\le\varepsilon,$$

from Proposition 2.1, this gives the relation valid for all $x_0 \ge 0$,

$$\mathbb{P}_{(x_0,x_1)}\left(\sup_{0\leq s\leq t}X_1^r(s)\geq D_1\right)\leq\varepsilon.$$

By Condition (C), there exists $\gamma \geq 0$ (that depends on x_1) such that $r(x_0, x_1) \geq \delta$ when $x_0 \geq \gamma$. As long as $(X^r(t))$ stays in the subset $\{(y_0, y_1) : y_1 \leq D_1\}$, the transition rates of the first component $(X_0^r(t))$ are uniformly bounded. Consequently, there exists K such that, for $x_0 \geq K$,

$$\mathbb{P}_{(x_0,x_1)}\left[\sup_{s\leq t} X_0^r(s) \leq \gamma, \ \sup_{s\leq t} X_1^r(s) \leq D_1, \right] \leq \varepsilon.$$

Relation (10) follows from the last two inequalities and the identity

$$\mathbb{P}_{(x_0,x_1)}(\sigma \wedge \tau_{\delta} \leq t) \leq \mathbb{P}_{(x_0,x_1)}\left(\sup_{s \leq t} X_0^r(s) \leq \gamma\right).$$

One returns to the proof of Inequality (9). By definition of the Q-matrix of the process $(X^r(t))$,

$$\mathbb{E}_{(x_0,x_1)}(|X^r(t|) - |x|) = \lambda t - \nu \int_0^t \mathbb{E}_{(x_0,x_1)}(X_1^r(u))du, \ x \in \mathbb{N}^2, \ t \ge 0.$$

For any $x \in \mathbb{N}^2$, there exists a version of $(Y^{\delta}(t))$ with initial condition $Y^{\delta}(0) =$ $X^{r}(0) = x$, and such that Relation (6) holds for $t < \tau_{\delta} \wedge \sigma$, in particular

$$\mathbb{E}_x(X_1^r(t)) \ge \mathbb{E}_x(X_1^r(t); t < \tau_{\delta} \land \sigma) \\ \ge \mathbb{E}_x(Y_1^{\delta}(t); t < \tau_{\delta} \land \sigma) = \mathbb{E}_x(Y_1^{\delta}(t)) - \mathbb{E}_x(Y_1^{\delta}(t); t \ge \tau_{\delta} \land \sigma).$$

Cauchy-Schwarz inequality shows that for any $t \ge 0$ and $x \in \mathbb{N}^2$

$$\int_{0}^{t} \mathbb{E}_{x}(Y_{1}^{\delta}(u);\tau_{\delta}\wedge\sigma\leq u) \, du \leq \int_{0}^{t} \sqrt{\mathbb{E}_{x}\left[\left(Y_{1}^{\delta}(u)\right)^{2}\right]} \sqrt{\mathbb{P}_{x}(\tau_{\delta}\wedge\sigma\leq u)} \, du$$
$$\leq \sqrt{\mathbb{P}_{x}(\tau_{\delta}\wedge\sigma\leq t)} \int_{0}^{t} \sqrt{\mathbb{E}_{x}\left[\left(Y_{1}^{\delta}(u)\right)^{2}\right]} \, du,$$

by gathering these inequalities, and by using the fact that the process $(Y_1^{\delta}(t))$ depends only on x_1 and not x_0 , one finally gets the relation

(11)
$$\frac{1}{t}\mathbb{E}_x(|X(t)| - |x|) \le \lambda - \frac{\nu}{t} \int_0^t \mathbb{E}_{x_1}(Y_1^{\delta}(u)) \, du + c(x_1, t)\sqrt{\mathbb{P}_x(\tau_{\delta} \land \sigma \le t)}$$

with

$$c(x_1,t) = \frac{\nu}{t} \int_0^t \sqrt{\mathbb{E}_{x_1}\left[\left(Y_1^{\delta}(u)\right)^2\right]} du.$$

Two cases are considered.

(1) If $\mu > \nu$, if $\delta < 1$ is such that $\delta \mu > \nu$, the process $(Y_1^{\delta}(t))$ is transient, so that

$$\lim_{t \to +\infty} \frac{1}{t} \int_0^t \mathbb{E}_{x_1}(Y_1^{\delta}(u)) \, du = +\infty,$$

for each $x_1 \ge 0$.

(2) If $\mu < \nu$, one takes $\delta = 1$, or if $\mu = \nu$, one takes $\delta < 1$ close enough to 1 so that $\lambda < \lambda^*(\delta)$. In both cases, $\lambda < \lambda^*(\delta)$ and the process $(Y_1^{\delta}(t))$ converges in distribution, hence

$$\lim_{t \to +\infty} \frac{1}{t} \int_0^t \mathbb{E}_{x_1}(Y_1^{\delta}(u)) \, du = \nu \mathbb{E}\left(Y_1^{\delta}(\infty)\right) = \lambda^*(\delta) > \lambda$$

for each $x_1 \ge 0$.

Consequently in both cases, there exist constants $\eta > 0$, $\delta < 1$ and $t_0 > 0$ such that for any $x_1 \leq K_1$,

(12)
$$\lambda - \nu \frac{1}{t_0} \int_0^{t_0} \mathbb{E}_{x_1}(Y_1^{\delta}(u)) du \le -\eta,$$

with Relation (11), one gets that if $x_1 \leq K_1$ then

$$\frac{1}{t_0}\mathbb{E}_x(|X(t_0)| - |x|) \le -\eta + c^*\sqrt{\mathbb{P}_x(\tau_\delta \wedge \sigma \le t_0)},$$

where $c^* = \max(c(n, t_0), 0 \le n \le K_1)$. By Identity (10), there exists K_0 such that, for all $x_0 \ge K_0$ and $x_1 \le K_1$, the relation

$$c^* \sqrt{\mathbb{P}_{(x_0,x_1)}(\tau_\delta \wedge \sigma \le t_0)} \le \frac{\eta}{2}$$

holds. This relation and the inequalities (12) and (11) give Inequality (9). The proposition is proved.

Another Boundary Condition. The boundary condition $x_1 \vee 1$ in the transition rates of (X(t)), Equation (2), prevents the second coordinate from ending up in the absorbing state 0. It amounts to suppose that a permanent server gets activated when no node may offer the file. Another way to avoid this absorbing state is to suppose that a permanent node is always active, which gives transition rates with $x_1 + 1$ instead. This choice was for instance made in Núñez-Queija and Prabhu [NQP08]. All our results apply for this other boundary condition: the only difference that is when $\nu > \mu$, the value of the threshold λ^* of Equation (4) is given by the quantity $\lambda^* = \mu\nu/(\nu - \mu)$.

3. Yule Processes with Deletions

This section introduces the tools which are necessary in order to generalize the results of the previous section to the multi-chunk case $n \ge 2$. A Yule process (Y(t)) with rate $\mu > 0$ is a Markovian branching process with Q-matrix

(13)
$$q_Y(x, x+1) = \mu x, \quad \forall x \ge 0.$$

An individual gives birth to a child, or equivalently splits into two particles, with rate μ . Let (σ_n) be the split times of a Yule process started with one particle, it is not difficult to check that, for $n \ge 1$,

$$\sigma_n \stackrel{\text{dist.}}{=} \sum_{\ell=1}^n \frac{E_{\ell}^{\mu}}{\ell} \stackrel{\text{dist.}}{=} \max(E_1^{\mu}, \dots, E_n^{\mu}),$$

where (E_{ℓ}^{μ}) are i.i.d. exponential random variables with parameter μ . If $\lambda > \mu$ then, by using Fubini's Theorem,

$$\mathbb{E}\left(\sum_{\ell=1}^{+\infty} e^{-\lambda\sigma_{\ell}}\right) = \mathbb{E}\left(\sum_{\ell=1}^{+\infty} \int_{0}^{+\infty} \lambda e^{-\lambda x} \mathbb{1}_{\{\sigma_{\ell} \le x\}} dx\right) = \int_{0}^{+\infty} \lambda e^{-\lambda x} \sum_{\ell=1}^{+\infty} \mathbb{P}(\sigma_{\ell} \le x) dx$$

$$(14) \qquad \qquad = \int_{0}^{+\infty} \lambda e^{-\lambda x} \frac{1 - e^{-\mu x}}{e^{-\mu x}} dx = \frac{\mu}{\lambda - \mu} < +\infty.$$

In this section one considers some specific results on variants of this stochastic model when some individuals are killed. In terms of branching processes, this amounts to prune the tree, i.e., to cut some edges of the tree, and the subtree attached to it. This procedure is fairly common for branching processes, in the Crump-Mode-Jagers model for example, see Kingman [Kin75]. See also Neveu [Nev86] or Aldous and Pitman [AP98]. Two situations are considered: the first one when the deletions are part of the internal dynamics, so that each individual dies out after an exponential time, and the other when killings are given by an exogenous process and occur at fixed (random or deterministic) epochs. Constant Death Rate and Regeneration. Let (Z(t)) be the birth-and-death process whose Q-matrix Q_Z is given by, for $\mu_Z > 0$ and $\nu > 0$,

(15)
$$q_Z(z, z+1) = \mu_Z(z \lor 1)$$
 and $q_Z(z, z-1) = \nu z$.

The lifetime of an individual is exponentially distributed with parameter ν , and the process restarts with one individual after some time when it hits 0. This process can be described equivalently as a time-changed M/M/1 queue or as a sequence of independent branching processes. As it will be seen these two viewpoints are complementary.

In the rest of this part, μ_Z and ν are fixed, (Z(t)) is the Markov process with *Q*-matrix Q_Z , (σ_n) is the sequence of times of its positive jumps, the birth instants, and $(B_{\sigma}(t))$ is the corresponding counting process of (σ_n) , for $t \ge 0$,

$$B_{\sigma}(t) = \sum_{i \ge 1} \mathbb{1}_{\{\sigma_i \le t\}}$$

PROPOSITION 3.1 (Queueing Representation). If $Z(0) = z \in \mathbb{N}$, then

(16)
$$(Z(t), t \ge 0) \stackrel{\text{dist.}}{=} (L(C(t)), t \ge 0),$$

where (L(t)) is the process of the number of jobs of an M/M/1 queue with input rate μ_Z and service rate ν and with L(0) = z and $C(t) = \inf \{s > 0 : A(s) > t\}$, where

$$A(t) = \int_0^t \frac{1}{1 \vee L(u)} \, du.$$

PROOF. It is not difficult to check that the process $(M(t)) \stackrel{\text{def.}}{=} (L(C(t)))$ has the Markov property. Let Q_M be its *Q*-matrix. For $z \ge 0$,

$$\mathbb{P}(L(C(h)) = z + 1 \mid L(0) = z) = \mu_Z \mathbb{E}(C(h)) + o(h) = \mu_Z(z \lor 1)h + o(h),$$

hence $q_M(z, z+1) = \mu_Z(z \vee 1)$. Similarly $q_M(z, z-1) = \nu z$. The proposition is proved.

COROLLARY 3.1. For any $\gamma > (\mu_Z - \nu) \lor 0$ and $z = Z(0) \in \mathbb{N}$,

(17)
$$\mathbb{E}_z\left(\sum_{n=1}^{+\infty} e^{-\gamma\sigma_n}\right) < +\infty.$$

PROOF. Proposition 3.1 shows that, in particular, the sequences of positive jumps of (Z(t)) and of (L(C(t))) have the same distribution. Hence, if $\mathcal{N}_{\mu_Z} = (t_n)$ is the arrival process of the M/M/1 queue, a Poisson process with parameter μ_Z , then, with the notations of the above proposition, the relation

$$(\sigma_n) \stackrel{\text{dist.}}{=} (A(t_n))$$

holds. By using standard martingale properties of stochastic integrals with respect to Poisson processes, see Rogers and Williams [**RW87**], one gets for $t \ge 0$,

$$\mathbb{E}_{z}\left(\sum_{n\geq 1}e^{-\gamma A(t_{n})}\right) = \mathbb{E}_{z}\left(\int_{0}^{\infty}e^{-\gamma A(s)}\mathcal{N}_{\mu_{Z}}(ds)\right) = \mu_{Z}\mathbb{E}_{z}\left(\int_{0}^{\infty}e^{-\gamma A(s)}\,ds\right)$$
(18)
$$= \mu_{Z}\int_{0}^{\infty}e^{-\gamma u}\mathbb{E}_{z}\left(Z(u)\vee 1\right)\,du,$$

where Relation (16) has been used for the last equality. Kolmogorov's equation for the process (Z(t)) gives that

$$\phi(t) \stackrel{\text{def.}}{=} \mathbb{E}_z(Z(t)) = \mu_Z \int_0^t \mathbb{E}_z \left(Z(u) \lor 1 \right) \, du - \nu \int_0^t \mathbb{E}_z \left(Z(u) \right) \, du$$
$$\leq \left(\mu_Z - \nu \right) \int_0^t \phi(u) \, du + \mu_Z t,$$

therefore, by Gronwall's Lemma,

$$\phi(t) \le \phi(0) + \mu_Z \int_0^t u e^{(\mu_Z - \nu)u} \, du \le z + \frac{\mu_Z}{\mu_Z - \nu} t e^{(\mu_Z - \nu)t}.$$

From Equation (18), one concludes that

$$\mathbb{E}_{z}\left(\sum_{n}e^{-\gamma\sigma_{n}}\right) = \mathbb{E}_{z}\left(\sum_{n}e^{-\gamma A(t_{n})}\right) < +\infty.$$

The proposition is proved.

A Branching Process Before hitting 0, the Markov process (Z(t)) whose Q-matrix is given by Relation (15) can be seen a Bellman-Harris branching process. Its Malthusian parameter is given by $\alpha = \mu_Z - \nu$. See Athreya and Ney [AN72]. In this setting, it describes the evolution of a population of independent particles, at rate $\lambda \stackrel{\text{def.}}{=} \mu_Z + \nu$ each of these particles either splits into two particles with probability $p \stackrel{\text{def.}}{=} \mu_Z / (\mu_Z + \nu)$ or dies. These processes will be referred to as (p, λ) branching processes in the sequel.

A (p, λ) -branching process survives with positive probability only when p > 1/2, in which case the probability of extinction q is equal to $q = (1-p)/p = \nu/\mu_Z$. The main (and only) difference with a branching process is that Z regenerates after hitting 0. When it regenerates, it again behaves as a (p, λ) -branching process (started with one particle), until it hits 0 again.

PROPOSITION 3.2 (Branching Representation). If $Z(0) = z \in \mathbb{N}$ and $(\widetilde{Z}(t))$ is a (p, λ) -branching process started with $z \in \mathbb{N}$ particles and \widetilde{T} its extinction time, then

$$(Z(t), 0 \le t \le T) \stackrel{dist.}{=} (\widetilde{Z}(t), 0 \le t \le \widetilde{T}).$$

where $T = \inf\{t \ge 0 : Z(t) = 0\}$ is the hitting time of 0 by (Z(t)).

COROLLARY 3.2. Suppose that $\mu_Z > \nu$. Then \mathbb{P}_z -almost surely for any $z \ge 0$, there exists a finite random variable $Z(\infty)$ such that,

$$\lim_{t \to +\infty} e^{-(\mu_Z - \nu)t} Z(t) = Z(\infty) \quad and \quad Z(\infty) > 0.$$

PROOF. When $\mu_Z > \nu$, the process (Z(t)) couples in finite time with a supercritical (p, λ) -branching process $(\widetilde{Z}(t))$ conditioned on non-extinction; this follows readily from Proposition 3.2 (or see the Appendix for details). Since for any supercritical (p, λ) -branching process, $(\exp(-(\mu_Z - \nu)t)\widetilde{Z}(t))$ converges almost surely to a finite random variable $\widetilde{Z}(\infty)$, positive on the event of non-extinction (see Nerman [Ner81]), one gets the desired result.

Due to its technicality, the proof of the following result is postponed to the Appendix; this result is used in the proof of Proposition 3.5.

PROPOSITION 3.3. Suppose that $\mu_Z > \nu$, if

(19)
$$\eta^*(x) = \frac{2 - x - \sqrt{x(4 - 3x)}}{2(1 - x)}, \ 0 < x < 1.$$

then for any $0 < \eta < \eta^*(\nu/\mu_Z)$,

$$\sup_{z\geq 0} \left[\mathbb{E}_z \left(\sup_{t\geq \sigma_1} \left(e^{\eta(\mu_Z - \nu)t} B_{\sigma}(t)^{-\eta} \right) \right) \right] < +\infty.$$

A Yule Process Killed at Fixed Instants. In this part, it is assumed that, provided that it is non-empty, at epochs σ_n , $n \geq 1$, an individual is removed from the population of an ordinary Yule process (Y(t)) with rate μ_W starting with $Y(0) = w \in \mathbb{N}$ individuals. It is assumed that (σ_n) is some fixed non-decreasing sequence. It will be shown that the process (W(t)) obtained by killing one individual of (Y(t)) at each of the successive instants (σ_n) survives with positive probability when the series with general term $(\exp(-\mu_W \sigma_n))$ converges.

In the following, a related result will be considered in the case where the sequence (σ_n) is given by the sequence of birth times of the process (Z(t)) introduced above. See Alsmeyer [Als93] and the references therein for related models.

One denotes

$$\kappa = \inf\{n \ge 1 : W(\sigma_n) = 0\}$$

The process (W(t)) can be represented in the following way

(20)
$$W(t) = Y(t) - \sum_{i=1}^{\kappa} X_i(t) \mathbb{1}_{\{\sigma_i \le t\}},$$

where, for $1 \leq i \leq \kappa$ and $t \geq \sigma_i$, $X_i(t)$ is the total number of children at time t in the original Yule process of the *i*th individual killed at time σ_i . In terms of trees, (W(t)) can be seen as a subtree of (Y(t)): for $1 \le i \le \kappa$, $(X_i(t))$ is the subtree of (Y(t)) associated with the *i*th particle killed at time σ_i .

It is easily checked that $(X_i(t - \sigma_i), t \ge \sigma_i)$ is a Yule process starting with one individual and, since a killed individual cannot have one of his descendants killed, that the processes

$$(\widetilde{X}_i(t)) = (X_i(t+\sigma_i), t \ge 0), \quad 1 \le i \le \kappa,$$

are independent Yule processes.

For any process (U(t)), one denotes

(21)
$$(M_U(t)) \stackrel{\text{def.}}{=} \left(e^{-\mu_W t} U(t) \right).$$

If (X(t)) is a Yule process with rate μ_W , the martingale $(M_{\widetilde{X}}(t))$ converges almost surely and in L_2 to a random variable $M_{\widetilde{X}}(\infty)$ with an exponential distribution with mean $\widetilde{X}(0)$, and by Doob's Inequality

$$\mathbb{E}\left(\sup_{t\geq 0} M_{\widetilde{X}}(t)^2\right) \leq 2\sup_{t\geq 0} \mathbb{E}\left(M_{\widetilde{X}}(t)^2\right) < +\infty.$$

See Athreva and Nev [AN72]. Consequently

$$e^{-\mu_W t}W(t) = M_Y(t) - \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(t-\sigma_i) \mathbb{1}_{\{\sigma_i \le t\}},$$

and for any $t \ge 0$,

$$\sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(t-\sigma_i) \mathbb{1}_{\{\sigma_i \le t\}} \le \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} \sup_{s \ge 0} M_{\widetilde{X}_i}(s).$$

Assume now that $\sum_{i\geq 1} e^{-\mu_W \sigma_i} < +\infty$: then the last expression is integrable, and Lebesgue's Theorem implies that $(M_W(t)) = (\exp(-\mu_W t)W(t))$ converges almost surely and in L_2 to

$$M_W(\infty) = M_Y(\infty) - \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(\infty).$$

Clearly, for some w^* large enough and then for any $w \ge w^*$, one has

$$\mathbb{E}_w(M_W(\infty)) \ge w - \sum_{i=1}^{+\infty} e^{-\mu_W \sigma_i} > 0,$$

in particular $\mathbb{P}_w(M_W(\infty) > 0) > 0$ and $\mathbb{P}_w(W(t) \ge 1, \forall t \ge 0) > 0$. If $Y(0) = w < w^*$ and $\sigma_1 > 0$, then $\mathbb{P}_w(Y(\sigma_1) \ge w^* + 1) > 0$ and therefore, by translation at time σ_1 , the same conclusion holds when the sequence $(\exp(-\mu_W \sigma_i))$ has a finite sum. The following proposition has thus been proved.

PROPOSITION 3.4. Let (W(t)) be a process growing as a Yule process with rate μ_W and for which individuals are killed at non-decreasing instants (σ_n) with $\sigma_1 > 0$. If

$$\sum_{i=1}^{+\infty} e^{-\mu_W \sigma_i} < +\infty,$$

then as t gets large, and for any $w \ge 1$, the variable $(\exp(-\mu_W t)W(t))$ converges \mathbb{P}_w -almost surely and in L_2 to a finite random variable $M_W(\infty)$ such that $\mathbb{P}_w(M_W(\infty) > 0) > 0$.

The previous proposition establishes the minimal results needed in Section 4. However, Kolmogorov's Three-Series, see Williams [Wil91], can be used in conjunction with Fatou's Lemma to show that (W(t)) dies out almost surely when the series with general term $(\exp(-\mu_W \sigma_n))$ diverges.

A Yule Process Killed at the Birth Instants of a Bellman-Harris Process. In this subsection, one considers a Yule process (Y(t)) with parameter μ_W with Qmatrix defined by Relation (13) and an independent Markov process (Z(t)) with Q-matrix defined by Relation (15). In particular $\mu_Z - \nu$ is the Malthusian parameter of (Z(t)). A process (W(t)) is defined by killing one individual of (Y(t)) at each of the birth instants (σ_n) of (Z(t)). As before $(B_{\sigma}(t))$ denotes the counting process association to the non-decreasing sequence (σ_n) ,

$$B_{\sigma}(t) = \sum_{i \ge 1} \mathbb{1}_{\{\sigma_i \le t\}}.$$

PROPOSITION 3.5. Assume that $\mu_Z - \nu > \mu_W$, and let H_0 be the extinction time of (W(t)), i.e.,

$$H_0 = \inf\{t \ge 0 : W(t) = 0\}$$

then the random variable H_0 is almost surely finite and:

(i)
$$Z(H_0) - Z(0) \le e^{\mu_W H_0} M_V^*$$
 where

$$M_Y^* = \sup_{t \ge 0} e^{-\mu_W t} Y(t)$$

(ii) There exists a finite constant C such that for any $z \ge 0$ and $w \ge 1$,

(22)
$$\mathbb{E}_{(w,z)}(H_0) \le C\left(\log(w) + 1\right).$$

Note that the subscript (w, z) refers to the initial state of the Markov process (W(t), Z(t)).

PROOF. Define $\alpha = \mu_Z - \nu$. Concerning the almost sure finiteness of H_0 , note that Equation (20) entails that $W(t) \leq Y(t) - B_{\sigma}(t)$ for all $t \geq 0$ on the event $\{H_0 = +\infty\}$. As t goes to infinity, both $\exp(-\mu_W t)Y(t)$ and $\exp(-\alpha t)B_{\sigma}(t)$ converge almost surely to positive and finite random variables (see Nerman [Ner81]), which implies, when $\alpha = \mu_Z - \nu > \mu_W$, that W(t) converges to $-\infty$ on $\{H_0 = +\infty\}$, and so this event is necessarily of probability zero.

The first point (i) of the proposition comes from Identity (20) at $t = H_0$:

(23)
$$Z(H_0) - Z(0) \le B_{\sigma}(H_0) \le Y(H_0) \le e^{\mu_W H_0} M_Y^*.$$

By using the relation $\exp(x) \ge x$, Equation (22) follows from the following bound: for any $\eta < \eta^*(\nu/\mu_Z)$ (recall that η^* is given by Equation (19)),

(24)
$$\sup_{w \ge 1, z \ge 0} \left[w^{-\eta} \mathbb{E}_{(w,z)} \left(e^{\eta (\alpha - \mu_W) H_0} \right) \right] < +\infty.$$

So all is left to prove is this bound. Under $\mathbb{P}_{(w,z)}$, (Y(t)) can be represented as the sum of w i.i.d. Yule processes, and so $M_Y^* \leq M_{Y,1}^* + \cdots + M_{Y,w}^*$ with $(M_{Y,i}^*)$ i.i.d. distributed like M_Y^* under $\mathbb{P}_{(1,z)}$; Inequality (23) then entails that

$$e^{(\alpha-\mu_W)H_0} \le \left(\sum_{i=1}^w M_{Y,i}^*\right) \times \sup_{t\ge\sigma_1} \left(e^{\alpha t}/B_{\sigma}(t)\right).$$

By independence of $(M_{Y,i}^*)$ and $(B_{\sigma}(t))$, Jensen's inequality gives for any $\eta < 1$

$$\mathbb{E}_{(w,z)}\left(e^{\eta(\alpha-\mu_W)H_0}\right) \le w^{\eta}\left(\mathbb{E}\left(M_{Y,1}^*\right)\right)^{\eta}\mathbb{E}_z\left(\sup_{t\ge\sigma_1}\left(e^{\eta\alpha t}B_{\sigma}(t)^{-\eta}\right)\right)$$

hence the bound (24) follows from Proposition 3.3.

One concludes this section with a Markov chain which will be used in Section 4. Define recursively the sequence (V_n) by, $V_0 = v$ and

(25)
$$V_{n+1} = \sum_{k=1}^{A_n(V_n)} I_k, \ n \ge 0,$$

where (I_k) are identically distributed integer valued random variables independent of V_n and $A_n(V_n)$, and such that $\mathbb{E}(I_1) = p$ for some $p \in (0, 1)$. For v > 0, $A_n(v)$ is an independent random variable with the same distribution as $Z(H_0)$ under $\mathbb{P}_{(1,v)}$, i.e., with the initial condition (W(0), Z(0)) = (1, v).

The above equation (25) can be interpreted as a branching process with immigration, see Seneta [Sen70], or also as an autoregressive model.
PROPOSITION 3.6. Under the condition $\mu_Z - \nu > \mu_W$, if (V_n) is the Markov chain defined by Equation (25) and, for $K \ge 0$,

$$N_K = \inf\{n \ge 0 : V_n \le K\},\$$

then there exist $\gamma > 0$ and $K \in \mathbb{N}$ such that

(26)
$$\mathbb{E}(N_K|V_0=v) \le \frac{1}{\gamma}\log(1+v), \quad \forall v \ge 0.$$

The Markov chain (V_n) is in particular positive recurrent.

PROOF. For $V_0 = v \in \mathbb{N}$, Jensen's Inequality and Definition (25) give the relation

(27)
$$\mathbb{E}_{v} \log\left(\frac{1+V_{1}}{1+v}\right) \leq \mathbb{E}_{(1,v)} \log\left[\frac{1+pZ(H_{0})}{1+v}\right].$$

From Proposition 3.5 and by using the same notations, one gets that, under $\mathbb{P}_{(1,v)}$,

$$Z(H_0) \le v + e^{\mu_W H_0} M_Y^*$$

where (Y(t)) is a Yule process starting with one individual. By looking at the birth instants of (Z(t)), it is easily checked that the random variable H_0 under $\mathbb{P}_{(1,v)}$ is stochastically bounded by H_0 under $\mathbb{P}_{(1,0)}$. The integrability of H_0 under $\mathbb{P}_{(1,0)}$ (proved in Proposition 3.5) and of M_Y^* give that the expression

$$\log\left(\frac{1+p(v+e^{\mu_W H_0}M_Y^*)}{1+v}\right)$$

bounding the right hand side of Relation (27) is also an integrable random variable under $\mathbb{P}_{(1,0)}$. Lebesgue's Theorem gives therefore that

$$\limsup_{v \to +\infty} \left[\mathbb{E}_v \log \left(\frac{1+V_1}{1+v} \right) \right] \le \log p < 0$$

Consequently, one concludes that $v \mapsto \log(1+v)$ is a Lyapunov function for the Markov chain (V_n) , i.e., if $\gamma = -(\log p)/2$, there exists K such that for $v \ge K$,

$$\mathbb{E}_v \log \left(1 + V_1\right) - \log \left(1 + v\right) \le -\gamma.$$

Foster's criterion, see Theorem 8.6 of Robert [**Rob03**], implies that (V_n) is indeed ergodic and that Relation (26) holds.

4. Analysis of the Multi-Chunk Network

In this section it is assumed that a file of n chunks is distributed by the filesharing network within the following framework, corresponding to Figure 1. Chunks are delivered in the sequential order, and, for $k \ge 1$, requests with chunks $1, \ldots, k$ provide service for requests with one less chunk.

For $0 \le k < n$ and $t \ge 0$, the variable $X_k(t)$ denotes the number of requests downloading the (k+1)st chunk; for k = n, $X_n(t)$ is the number of requests having all the chunks. When taking into account the boundaries in the transition rates described in Figure 1, one gets the following Q-matrix for the (n+1)-dimensional Markov process $(X_k(t), 0 \le k \le n)$:

$$Q(f)(x) = \lambda [f(x+e_0) - f(x)] + \sum_{k=1}^{n} \mu_k (x_k \vee 1) [f(x+e_k - e_{k-1}) - f(x)] \mathbb{1}_{\{x_{k-1} > 0\}} + \nu x_n [f(x-e_n) - f(x)],$$

where $x \in \mathbb{N}^{n+1}$, $f : \mathbb{N}^{n+1} \to \mathbb{R}_+$ is a function and, for $0 \le k \le n$, $e_k \in \mathbb{N}^{n+1}$ is the kth unit vector. Note that, as before, to avoid absorbing states, it is assumed that there is a server for the kth chunk when $x_k = 0$. The first section corresponds to the case n = 1 in a more general setting.

It is first shown in Proposition 4.1 that the network is stable for sufficiently small input rate λ . Proposition 4.2 studies the analog of the two-dimensional case with $\mu > \nu$, i.e., when $\mu_1 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$, it is proved that the network is stable for any input rate λ . When this condition fails, it is shown that for n = 2 the network can only accommodate a finite input rate.

PROPOSITION 4.1. Under the condition

(28)
$$\sum_{k=1}^{n} \frac{\lambda}{\mu_k} < 1,$$

the Markov process (X(t)) is ergodic for any $\nu > 0$.

Condition (28) is obviously not sharp as can be seen in the case n = 1 analyzed in Section 2. But the proposition shows that there is always a positive threshold λ^* such that the system is stable when $\lambda < \lambda^*$.

PROOF. For $x \in \mathbb{N}^{n+1}$ and $(\alpha_k) \in \mathbb{R}^{n+1}$, define $f(x) = \alpha_0 x_0 + \dots + \alpha_n x_n$, then

$$Q(f)(x) = \lambda \alpha_0 - \sum_{k=1}^{n} (\alpha_{k-1} - \alpha_k) \mu_k(x_k \vee 1) \mathbb{1}_{\{x_{k-1} > 0\}} - \nu x_n \alpha_n.$$

For $\varepsilon > 0$, one can choose (α_k) so that $\alpha_0 = 1$ and

$$\alpha_{k-1} - \alpha_k = \frac{\lambda}{\mu_k} + \varepsilon, \quad 1 \le k \le n,$$

hence

$$\alpha_n = 1 - \left(n\varepsilon + \sum_{i=1}^n \frac{\lambda}{\mu_k} \right),$$

so that, for ε small enough, the α_k 's, $0 \le k \le n$ are decreasing and positive under the condition of the proposition; in particular the set $\{x : f(x) \le K\}$ is finite for any $K \ge 0$.

Take $K = (1 + \lambda)/\nu$, then if $x \in \mathbb{N}^{n+1}$ is such that $f(x) \ge K$, either $x_k > 0$ for some $0 \le k \le n-1$ and in this case

$$Q(f)(x) \le \lambda - \mu_{k+1}(\alpha_k - \alpha_{k+1}) = -\varepsilon \mu_{k+1} < 0,$$

or $x_n \geq K$ so that

$$Q(f)(x) \le \lambda - \nu K = -1 < 0.$$

A Lyapunov function criteria for Markov processes shows that this implies that the Markov process (X(t)) is ergodic. See Proposition 8.14 of Robert [**Rob03**] for example.

Decreasing Service Rates. The analog of the "good" two-dimensional case $\mu > \nu$ is proved in the next proposition.

PROPOSITION 4.2. Under the condition $\mu_1 > \mu_2 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$, the Markov process $(X(t)) = (X_k(t), 0 \le k \le n)$ describing the linear file-sharing network is ergodic for any $\lambda \ge 0$.

PROOF. The proof proceeds in two steps: first coupling arguments with Yule processes allow to prove (30); then one can use the same technique as in the proof of Proposition 2.3, see Robert [**Rob03**, Theorem 9.7].

Step 1 (coupling). Let $(W_n(t))$ be the process with Q-matrix defined by Relation (15) with $\mu_Z = \mu_n$ and starting at $W_n(0) = w_n \ge 1$. Since $\mu_n > \nu$, the process $(\exp(-(\mu_n - \nu)t)W_n(t))$ converges almost surely to a finite and positive random variable $M_{W_n}(\infty)$ by Corollary 3.2. Moreover, since $\mu_{n-1} > \mu_n - \nu > 0$, Corollary 3.1 entails that the birth instants (σ_ℓ^n) of this process are such that

$$\sum_{\ell \ge 1} e^{-\mu_{n-1}\sigma_\ell^n} < +\infty, \quad \text{almost surely.}$$

Let $(Y_{n-1}(t))$ be an independent Yule process with parameter μ_{n-1} with initial condition $Y_{n-1}(0) = w_{n-1} \ge 1$ and $(W_{n-1}(t))$ the resulting process when its individuals are killed at the instants (σ_{ℓ}^n) of births of $(W_n(t))$: the previous equation and Proposition 3.4 show that $(W_{n-1}(t))$ can survive forever with a positive probability.

Let $(Y_{n-2}(t))$ be an independent Yule process starting from $w_{n-2} \ge 1$ with parameter μ_{n-2} . Define $(W_{n-2}(t))$ the resulting process when the individuals of $(Y_{n-2}(t))$ are killed at the birth instants (σ_{ℓ}^{n-1}) of $(W_{n-1}(t))$. Since $\mu_{n-2} > \mu_{n-1}$, the birth instants $(\widetilde{\sigma}_{\ell}^{n-1})$ of $(Y_{n-1}(t))$ satisfy

$$\sum_{\ell=1}^{+\infty} e^{-\mu_{n-2}\widetilde{\sigma}_{\ell}^{n-1}} < +\infty$$

almost surely by Equation (14) (which still holds for a Yule process starting with more than one particle). Since the birth instants (σ_{ℓ}^{n-1}) of $(W_{n-1}(t))$ are a subsequence of $(\tilde{\sigma}_{\ell}^{n-1})$, the same relationship holds for (σ_{ℓ}^{n-1}) , and therefore, with a positive probability, the three processes $(e^{-(\mu_n-\nu)t}W_n(t))$, $(e^{-\mu_{n-1}t}W_{n-1}(t))$ and $(e^{-\mu_{n-2}t}W_{n-2}(t))$ converge simultaneously to positive and finite random variables $M_{W_n}(\infty)$, $M_{W_{n-1}}(\infty)$ and $M_{W_{n-2}}(\infty)$, respectively. This construction can be repeated inductively to give the existence of n processes $(W_k(t), k = 1, \ldots, n)$ such that (σ_{ℓ}^k) is the sequence of birth times of W_k , W_n is the birth-and-death process with Q-matrix (15), W_k for $1 \le k \le n-1$ is a Yule process with parameter μ_k killed at (σ_{ℓ}^{k+1}) , and the event $\mathcal{E} = \{M_{W_1}(\infty) > 0, \ldots, M_{W_n}(\infty) > 0\}$ has a positive probability. On this event, $W_k(t) \ge 1$ for all $t \ge 0$ and $1 \le k \le n-1$, and

$$\lim_{t \to +\infty} W_n(t) = +\infty.$$

For $0 \le k \le n-1$, one defines $(X_k^S(t)) = (X_{k,n-k}^S(t), \ldots, X_{k,n}^S(t))$, the *k*th saturated system, as the (k+1)-dimensional Markov process with generator

(29)
$$Q_{k}^{S}(f)(x) = \mu_{n-k}(x_{n-k} \vee 1)[f(x+e_{n-k}) - f(x)] + \sum_{\ell=1}^{k} \mu_{n-k+\ell}(x_{n-k+\ell} \vee 1)[f(x+e_{n-k+\ell} - e_{n-k+\ell-1}) - f(x)]\mathbb{1}_{\{x_{n-k+\ell-1} > 0\}} + \nu x_{n}[f(x-e_{n}) - f(x)],$$

where $x \in \mathbb{N}^{k+1}$ and $f : \mathbb{N}^{k+1} \to \mathbb{R}_+$ is an arbitrary function. Compared with the process $(X_{\ell}(t), 1 \leq \ell \leq n)$ with generator Q, it amounts to look at the k+1last queues $(X_{n-k}(t), \ldots, X_n(t))$ under the assumption that the queue n-k-1 is saturated, i.e., $X_{n-k-1}(t) \equiv +\infty$ for all $t \geq 0$.

Note that for any $0 \le k \le n-1$, the transition rates of the Markov processes $(W_{n-\ell}(t), 0 \le \ell \le k)$ and $(X^S_{k,n-\ell}(t), 0 \le \ell \le k)$ are identical as long as no coordinate hits 0; one thus concludes that, with positive probability, the relation

$$\lim_{t \to +\infty} X_{k,n}^S(t) = +\infty$$

holds when $X_{k,n-\ell}^S(0) \ge 1$, $\ell = 0, \ldots, k$. Consequently, since the set $(\mathbb{N} - \{0\})^{k+1}$ can be reached with positive probability from any initial state in \mathbb{N}^{k+1} by $(X_k^S(t))$, then

(30)
$$\lim_{t \to +\infty} \mathbb{E}(X_{k,n}^S(t)) = +\infty.$$

Step 2 (Foster's criterion). We use Foster's criterion as stated in Theorem 9.7 of Robert [**Rob03**]. First we inspect the case when $X_n(0)$ is large, then the case when $X_n(0)$ is bounded and $X_{n-1}(0)$ is large, etc... The key idea is that when $X_{n-k-1}(0)$ is large, then the process $(X_{n-k}(t), \ldots, X_n(t))$ essentially behaves as the process $(X_k^S(t))$, for which Relation (30) ensures that the output rate is arbitrarily large.

Let $X(0) = x = (x_k) \in \mathbb{N}^{n+1}$, since the last queue serves at rate ν each request, for $t \ge 0$,

$$\mathbb{E}(\|X(t)\|) \le \|x\| + \lambda t - x_n \left(1 - e^{-\nu t}\right),$$

where $||x|| = x_0 + \cdots + x_n$ for $x = (x_0, \ldots, x_n) \in \mathbb{N}^{n+1}$. Define $t_n = 1$ and let K_n be such that $\lambda t_n - K_1(1 - \exp(-\nu)) \leq -1$, so that the relation

$$\mathbb{E}_x(||X(t_n)||) - ||x|| \le -1,$$

holds when $x_n \ge K_n$.

From Equation (30) with k = 0, one gets that there exists some t_{n-1} such that for any $x_n \leq K_n$,

$$\nu \int_{0}^{t_{n-1}} \mathbb{E}_{x_n} \left(X_{0,n}^S(u) \right) \, du \ge \lambda t_{n-1} + 2.$$

The two processes $(X_0^S(t))$ and (X(t)) can be built on the same probability space such that if they start from the same initial state, then the two processes $(X_{0,n}^S(t))$ and $(X_n(t))$ are identical as long as $X_{n-1}(t)$ stays positive. Since moreover the hitting time $\inf\{t \ge 0 : X_{n-1}(t) = 0\}$ goes to infinity as x_{n-1} goes to infinity for any $x_n \leq K_n$, one gets that there exists K_{n-1} such that if $x_{n-1} \geq K_{n-1}$ and $x_n < K_n$, then the relation

$$\mathbb{E}_{x}(\|X(t_{n-1})\|) - \|x\| = \lambda t_{n-1} - \nu \int_{0}^{t_{n-1}} \mathbb{E}_{x}(X_{n}(u)) du$$
$$\leq \lambda t_{n-1} - \left(\nu \int_{0}^{t_{n-1}} \mathbb{E}_{x_{n}}\left(X_{0,n}^{S}(u)\right) du - 1\right) \leq -1$$

holds.

By induction, one gets in a similar way that there exist constants t_n, \ldots, t_0 and K_n, \ldots, K_0 such that for any $0 \le \ell \le n$, if $x_n \le K_n, x_{n-1} \le K_{n-1}, \ldots, x_{n-\ell+1} \le K_{n-\ell+1}$ and $x_{n-\ell} > K_{n-\ell}$, then

$$\mathbb{E}_{x}(\|X(t_{n-\ell})\|) - \|x\| \le -1.$$

Theorem 8.13 of Robert [**Rob03**] shows that (X(t)) is an ergodic Markov process. The proposition is proved.

Analysis of the Two-Chunk Network. In this subsection, one investigates the case when the monotonicity condition $\mu_1 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$ fails. In general we conjecture the existence of bottlenecks which implies that the network can only accommodate a finite input rate. For instance, when $\mu_n - \nu < 0$, then it is easily seen that the network is unstable for $\lambda > \lambda^*$ where λ^* is defined in Equation (32) below.

The first non-trivial case occurs for n = 2, for which the monotonicity condition breaks in two situations, either when $\mu_2 - \nu > \mu_1$ or when $\mu_2 < \nu$. The latter case can be dealt in fact with the exact same arguments as before. See Proposition 4.4.

The actual difficulty is when $\mu_2 - \nu > \mu_1$: then the stationary behavior of $(X_2(t))$ is linked to the stationary behavior of the first saturated model $(X_1^S(t))$ defined through its *Q*-matrix (29). The difficulty in this case is that one needs to compare two processes which grow exponentially fast.

PROPOSITION 4.3. Assume that $\mu_2 - \nu > \mu_1$, then the first saturated process $(X_1^S(t))$ with Q-matrix defined by Equation (29) is ergodic.

COROLLARY 4.1. If $\mu_2 - \nu > \mu_1$ and if

$$\lambda_2^* \stackrel{aej.}{=} \nu \mathbb{E}_{\pi^S} \left(X_{1,2}^S(0) \right),$$

where π^S is the invariant distribution of the Markov process $(X_1^S(t))$, then the process $(X(t)) = (X_k(t), k = 0, 1, 2)$ describing the linear file-sharing network with parameters λ , μ_1 , μ_2 and ν is ergodic for $\lambda < \lambda_2^*$ and transient for $\lambda > \lambda_2^*$.

SKETCH OF PROOF. The proof of the transience when $\lambda > \lambda_2^*$ follows similarly as in Section 2: when $X_0(0)$ is large, the process $(X_1(t), X_2(t))$ can be coupled for some time with the second saturated system $(X_1^S(t))$. Since the output rate λ_2^* of this system is smaller than the input rate λ , this implies that $(X_0(t))$ builds up, and it can indeed be shown that $X_0(t)/t$ converges almost surely to $\lambda - \lambda_2^*$.

The ergodicity when $\lambda < \lambda_2^*$ is slightly more complicated, but it involves the same arguments as the ones employed in the proof of Proposition 4.2. The details are omitted.

PROOF OF PROPOSITION 4.3. Denote $(X_1^S(t)) = (X_{1,1}^S(t), X_{1,2}^S(t))$, then as long as the first coordinate $X_{1,1}^S$ is positive, the process $(X_1^S(t))$ has the same distribution as (W(t), Z(t)) introduced in Section 3: (Z(t)) is a Bellman-Harris process with Malthusian parameter $\mu_2 - \nu$ and (W(t)) is a Yule process with parameter μ_1 killed at times of births of (Z(t)).

By Proposition 3.5 and since $\mu_2 - \nu > \mu_1$, one has that $(X_{1,1}^S(t))$ returns infinitely often to 0. When $(X_{1,1}^S(t))$ is at 0 it jumps to 1 after an exponential time with parameter μ_1 , one denotes by $(E_{\mu_1,n})$ the corresponding i.i.d. sequence of successive residence times at 0. One defines the sequence (S_n) by induction, $S_0 = 0$ and then

$$S_{n+1} = \inf\{t > S_n : X_{1,1}^S(t) = 0\} + E_{\mu_1, n+1}, \ n \ge 0.$$

For $n \geq 1$, $X_{1,1}^S(S_n) = 1$ and for $n \geq 0$, define $M_n \stackrel{\text{def.}}{=} X_{1,2}^S(S_n)$. With the notations of Proposition 3.5, $(X_{1,1}^S(t))$ hits 0 after a duration of $H_{0,n}$ and at that time $(X_{1,2}^S(t))$ is at $Z(H_{0,n})$ with the initial condition $Z(0) = M_n$; while $X_{1,1}^S$ is still at 0, the dynamics of $X_{1,2}^S$ is simple, since it just empties. Finally, at time $S_{n+1} = S_n + H_{0,n} + E_{\mu_1,n+1}, (X_{1,1}^S(t))$ returns to 1 and at this instant the location of $(X_{1,2}^S(t))$ is given by

$$X_{1,2}^S(S_{n+1}) = M_{n+1} = \sum_{i=1}^{Z(H_{0,n})} \mathbb{1}_{\{E_{\nu,i} > E_{\mu_1,n+1}\}}$$

where $(E_{\nu,i})$ are i.i.d. exponential random variables with parameter ν , the *i*th variable being the residence time of the *i*th request in node 2. Consequently, $(M_n, n \ge 1)$ is a Markov chain whose transitions are defined by Relation (25) with $p = \nu/(\nu + \mu_1)$; note that $(M_n, n \ge 0)$ has the same dynamics only when $X_{1,1}^S(0) = 1$.

Define for any K > 0 the stopping time T_K

$$T_K = \inf\{t \ge 0 : X_{1,2}^S(t) \le K, X_{1,1}^S(t) = 1\}.$$

The ergodicity of $(X_1^S(t))$ will follow from the finiteness of $\mathbb{E}_{(x_1,x_2)}(T_K)$ for some K large enough and for arbitrary $x = (x_1, x_2) \in \mathbb{N}^2$. The strong Markov property of $(X_1^S(t))$ applied at time S_1 gives

$$\mathbb{E}_{(x_1,x_2)}(T_K) \le 2\mathbb{E}_{(x_1,x_2)}(S_1) + \mathbb{E}_{(x_1,x_2)} \left[\mathbb{E}_{(1,X_{1,2}^S(S_1))}(T_K) \right].$$

and so one only needs to study T_K conditioned on $\{X_{1,1}^S(0) = 1\}$ since $\mathbb{E}_{(x_1,x_2)}(S_1)$ is finite in view of Proposition 3.5.

Then, on this event and with N_K defined in Proposition 3.6, the identity

(31)
$$T_K = \sum_{i=0}^{N_K} (H_{0,i} + E_{\mu_1,i})$$

holds. For $i \ge 0$, the Markov property of $(M_n, n \ge 0)$ gives

$$\mathbb{E}_{(x_1,x_2)}\left(H_{0,i}\mathbb{1}_{\{i \le N_K\}}\right) = \mathbb{E}_{(x_1,x_2)}\left(\mathbb{E}_{(1,M_i)}\left(H_0\right)\mathbb{1}_{\{i \le N_K\}}\right)$$

With the same argument as in the proof of Proposition 3.6, one has

$$\mathbb{E}_{(1,M_i)}(H_0) \le \mathbb{E}_{(1,0)}(H_0) < +\infty,$$

with Equations (31) and (26) of Proposition (3.6), one gets that for some $\gamma > 0$ and some K > 0,

$$\mathbb{E}_{(x_1,x_2)}(T_K) \le 2\mathbb{E}_{(x_1,x_2)}(S_1) + C\left(1 + \mathbb{E}_{(x_1,x_2)}\left[\log\left(1 + X_{1,2}^S(S_1)\right)\right]\right)$$

with the constant $C = (\mathbb{E}_{(1,0)}(H_0) + 1/\mu_2)/\gamma$. This last term is finite for any (x_1, x_2) in view of Proposition 3.5, which proves the proposition.

PROPOSITION 4.4. If $\nu > \mu_2$ and

(32)
$$\lambda^* \stackrel{def.}{=} \frac{\mu_2}{(1 - \mu_2/\nu)(1 - \log(1 - \mu_2/\nu))},$$

then the Markov process $(X(t)) = (X_k(t), k = 0, 1, 2)$ is transient if $\lambda > \lambda^*$ and ergodic if $\lambda < \lambda^*$.

Sketch of Proof. The result for transience comes directly from the fact that the last coordinate is stochastically dominated by the birth-and-death process $(Y_1^1(t))$ of Section 2.

As before, the arguments employed in the proof of Proposition 4.2 to prove ergodicity can also be used, for this reason they are only sketched. One has in fact to consider the following situations.

- If there are many customers in the last queue, then the total number of customers instantaneously decreases.
- If there are many customers in the second queue, then the last queue has time to get close to stationarity, the input rate is λ and the output rate is λ*.
- Finally, if there are many customers in the first queue, then it is easily seen that the second queue builds up, since it grows like a Yule process killed at times (σ_n) where the sequence (σ_n) essentially grows linearly since the last queue is stable. Hence the second queue reaches high values and the last queue offers an output rate of λ^* .

Hence when $\lambda < \lambda^*$, the Markov process (X(t)) is ergodic.

Appendix A. Proof of Proposition 3.3

In this appendix the notations of Section 3 are used. Since the random variable $(B_{\sigma}(t) \mid Z(0) = 0)$ is stochastically smaller than $(B_{\sigma}(t) \mid Z(0) = z)$ for any $z \in \mathbb{N}$, it is enough to show that for $\eta < \eta^*(\nu/\mu_Z)$

$$\mathbb{E}_0\left[\sup_{t\geq\sigma_1}\left(e^{\eta\alpha t}B_{\sigma}(t)^{-\eta}\right)\right]<+\infty,$$

where $\alpha = \mu_Z - \nu > 0$.

Note that the process $(B_{\sigma}(t+\sigma_1), t \ge 0)$ under \mathbb{P}_0 has the same distribution as $(B_{\sigma}(t)+1, t \ge 0)$ under \mathbb{P}_1 , and by independence of σ_1 , an exponentially random variable with parameter μ_Z , and $(B_{\sigma}(t+\sigma_1), t \ge 0)$, one gets

$$\mathbb{E}_0\left[\sup_{t\geq\sigma_1}\left(e^{\eta\alpha t}B_{\sigma}(t)^{-\eta}\right)\right] = \mathbb{E}_0\left(e^{\eta\alpha\sigma_1}\right)\mathbb{E}_1\left[\sup_{t\geq0}\left(e^{\eta\alpha t}\left(B_{\sigma}(t)+1\right)^{-\eta}\right)\right].$$

Since $\alpha < \mu_Z$ and $\eta^*(\nu/\mu_Z) < 1$, then $\mathbb{E}_0(\exp(\eta\alpha\sigma_1))$ is finite, and all one needs to prove is that the second term is finite as well.

Define τ as the last time Z(t) = 0:

$$\tau = \sup\{t \ge 0 : Z(t) = 0\}$$

with the convention that $\tau = +\infty$ if (Z(t)) never returns to 0. Recall that, because of the assumption $\mu_Z > \nu$, with probability 1, the process (Z(t)) returns to 0 a finite number of times.

Conditioned on the event $\{\tau = +\infty\}$, the process (Z(t)) is a (p, λ) -branching process conditioned on survival, with $\lambda = \mu_Z + \nu$ and $p = \mu_Z/\lambda$. Such a branching process conditioned on survival can be decomposed as $Z = Z_{(1)} + Y$, where (Y(t)) is a Yule process (Y(t)) with parameter α . See Athreya and Ney [AN72]. Consequently, for any $0 < \eta < 1$,

$$\mathbb{E}_1\left[\sup_{t\geq 0} \left(e^{\eta\alpha t} \left(B_{\sigma}(t)+1\right)^{-\eta}\right) | \tau = +\infty\right] \leq \mathbb{E}_1\left[\sup_{t\geq 0} \left(e^{\eta\alpha t} Y(t)^{-\eta}\right)\right].$$

Since the *n*th split time t_n of (Y(t)) is distributed like the maximum of *n* i.i.d. exponential random variables, Y(t) for $t \ge 0$ is geometrically distributed with parameter $1 - e^{-\alpha t}$, hence,

$$\sup_{t\geq 0} \left[e^{\eta\alpha t} \mathbb{E}_1\left(\frac{1}{Y(t)^{\eta}}\right) \right] = \sup_{t\geq 0} \left[e^{-(1-\eta)\alpha t} \sum_{k\geq 1} \frac{(1-e^{-\alpha t})^{k-1}}{k^{\eta}} \right]$$
$$\leq \sup_{0\leq u\leq 1} \left[(1-u)^{1-\eta} \sum_{k\geq 1} \frac{u^{k-1}}{k^{\eta}} \right].$$

For 0 < u < 1, the relation

$$(1-u)^{1-\eta} \sum_{k\geq 1} \frac{u^{k-1}}{k^{\eta}} \leq (1-u)^{1-\eta} \int_0^\infty \frac{u^x}{(1+x)^{\eta}} \, dx,$$
$$= \left(\frac{1-u}{-\log u}\right)^{1-\eta} \int_0^\infty \frac{e^{-x}}{(x-\log u)^{\eta}} \, dx,$$

holds, hence

$$\sup_{t\geq 0} \left[e^{\eta \alpha t} \mathbb{E}_1\left(\frac{1}{Y(t)^\eta}\right) \right] < +\infty$$

The process $(e^{-\alpha t}Y(t))$ being a martingale, by convexity the process $(e^{\eta \alpha t}Y(t)^{-\eta})$ is a non-negative sub-martingale. For any $\eta \in (0,1)$ Doob's L_p inequality gives the existence of a finite $q(\eta) > 0$ such that

$$\mathbb{E}_1\left[\sup_{t\geq 0} \left(e^{\eta\alpha t}Y(t)^{-\eta}\right)\right] \leq q(\eta) \sup_{t\geq 0} \left[e^{\eta\alpha t}\mathbb{E}_1\left(\frac{1}{Y(t)^{\eta}}\right)\right] < +\infty.$$

The following result has therefore been proved.

LEMMA A.1. For any $0 < \eta < 1$,

$$\mathbb{E}_1\left[\sup_{t\geq 0} \left(e^{\eta\alpha t} \left(B_{\sigma}(t)+1\right)^{-\eta}\right) \middle| \tau = +\infty\right] < +\infty.$$

On the event $\{\tau < +\infty\}$, (Z(t)) hits a geometric number of times 0 and then couples with a (p, λ) -branching process conditioned on survival. On this event,

$$\sup_{t\geq 0} \left(e^{\eta\alpha t} \left(B_{\sigma}(t) + 1 \right)^{-\eta} \right)$$

= $\max \left(\sup_{0\leq t\leq \tau} \left(e^{\eta\alpha t} \left(B_{\sigma}(t) + 1 \right)^{-\eta} \right), \sup_{t\geq \tau} \left(e^{\eta\alpha t} \left(B_{\sigma}(t) + 1 \right)^{-\eta} \right) \right)$
 $\leq e^{\eta\alpha \tau} \left(1 + \sup_{t\geq 0} \left(e^{\eta\alpha t} \left(B'_{\sigma}(t) + 1 \right)^{-\eta} \right) \right)$

where $B'_{\sigma}(t)$ for $t \geq \tau$ is the number of births in $(\tau, t]$ of a (p, λ) -branching process conditioned on survival and independent of the variable τ , consequently

$$\mathbb{E}_{1}\left[\sup_{t\geq0}\left(e^{\eta\alpha t}\left(B_{\sigma}(t)+1\right)^{-\eta}\right)\middle|\tau<+\infty\right]\leq\mathbb{E}_{1}\left(e^{\eta\alpha\tau}\middle|\tau<+\infty\right)$$
$$\times\left(1+\mathbb{E}_{1}\left[\sup_{t\geq0}\left(e^{\eta\alpha t}\left(B_{\sigma}(t)+1\right)^{-\eta}\right)\middle|\tau=+\infty\right]\right).$$

In view of Lemma A.1, the proof of Proposition 3.3 will be finished if one can prove that

$$\mathbb{E}_1\left(e^{\eta\alpha\tau} \mid \tau < +\infty\right) < +\infty,$$

which actually comes from the following decomposition: under $\mathbb{P}_1(\cdot | \tau < +\infty)$, the random variable τ can be written as

$$\tau = \sum_{k=1}^{1+G} (T_k + E_{\mu_Z,k})$$

where G is a geometric random variable with parameter $q = \nu/\mu_Z$, (T_k) is an i.i.d. sequence with the same distribution as the extinction time of a (p, λ) -branching process starting with one particle and conditioned on extinction and $(E_{\mu_Z,k})$ are i.i.d. exponential random variables with parameter μ_Z .

Since q is the probability of extinction of a (p, λ) -branching process started with one particle, G + 1 represents the number of times (Z(t)) hits 0 before going to infinity. This representation entails

$$\mathbb{E}_1\left(e^{\eta\alpha\tau} \,|\, \tau < +\infty\right) = \mathbb{E}\left(\gamma(\eta)^{G+1}\right) \quad \text{where} \quad \gamma(\eta) = \mathbb{E}\left(e^{\eta\alpha(T_1 + E_{\mu_Z,1})}\right).$$

A (p, λ) -branching process conditioned on extinction is actually a $(1 - p, \lambda)$ branching process. See again Athreya and Ney [AN72]. Thus T_1 satisfies the following recursive distributional equation:

$$T_1 \stackrel{\text{dist.}}{=} E_{\lambda} + \mathbb{1}_{\{\xi=2\}} (T_1 \lor T_2),$$

where $\mathbb{P}(\xi = 2) = 1 - p$ and E_{λ} is an exponential random variable with parameter λ . This equation yields

$$\mathbb{P}(T_1 \ge t) \le e^{-\lambda t} + 2\lambda(1-p) \int_0^t \mathbb{P}(T_1 \ge t-u) e^{-\lambda u} \, du$$

and Gronwall's Lemma applied to the function $t \mapsto \exp(\lambda t) \mathbb{P}(T_1 \ge t)$ gives that

$$\mathbb{P}(T_1 \ge t) \le e^{(\lambda - 2\lambda p)t} = e^{(\nu - \mu_Z)t}$$

hence for any $0 < \eta < 1$,

$$\mathbb{E}_1(e^{\eta \alpha T_1}) \le \frac{1}{1-\eta}$$

Since G is a geometric random variable with parameter q, $\mathbb{E}(\gamma(\eta)^G)$ is finite if and only if $\gamma(\eta) < q$. Since finally

$$\gamma(\eta) = \frac{\mu_Z}{\mu_Z - \eta\alpha} \mathbb{E}\left(e^{\eta\alpha T_1}\right) \le \frac{\mu_Z}{(1 - \eta)(\mu_Z - \eta\alpha)},$$

one can easily check that $\gamma(\eta) < q$ for $\eta < \eta^*(\nu/\mu_Z)$ as defined by Equation (19), which concludes the proof of Proposition 3.3.

References of Chapter II

- [Als93] Gerold Alsmeyer. On the Galton-Watson Predator-Prey Process. Annals of Applied Probability, 3(1):198–211, 1993.
- [AN72] K. B. Athreya and P. E. Ney. *Branching processes*. Springer, 1972.
- [AP98] David Aldous and Jim Pitman. Tree-valued Markov chains derived from Galton-Watson processes. Ann. Inst. H. Poincaré Probab. Statist., 34(5):637–686, 1998.
- [BMM⁺08] Thomas Bonald, Laurent Massoulié, Fabien Mathieu, Diego Perino, and Andrew Twigg. Epidemic live streaming: optimal performance trade-offs. In *Proceedings of SIGMETRICS'08*, pages 325–336, New York, NY, USA, 2008. ACM.
- [Bra08] Maury Bramson. Stability of queueing networks, volume 1950 of Lecture Notes in Mathematics. Springer, Berlin, 2008. Lectures from the 36th Probability Summer School held in Saint-Flour, July 2–15, 2006.
- [CY01] Hong Chen and David D. Yao. Fundamentals of queueing networks. Springer-Verlag, New York, 2001. Performance, asymptotics, and optimization, Stochastic Modelling and Applied Probability.
- [DPM07] Trang Dinh Dang, R. Pereczes, and S. Molnar. Modeling the population of file-sharing peer-to-peer networks with branching processes. *Computers and Communications, 2007. ISCC 2007. 12th IEEE Symposium on*, pages 809–815, 2007.
- [Kel91] F. P. Kelly. Loss networks. The Annals of Applied Probability, 1(3):319– 378, 1991.
- [Kin75] J. F. C. Kingman. The first birth problem for an age-dependent branching process. *Annals of Probability*, 3(5):790–801, October 1975.
- [Les09] Lasse Leskelä. Stochastic Relations of Random Variables and Processes. *Journal of Theoretical Probability*, to appear, 2009.
- [MT08] Laurent Massoulié and Andrew Twigg. Rate-optimal schemes for Peerto-Peer live streaming. *Performance Evaluations*, 65(11-12):804–822, 2008.
- [MV05] Laurent Massoulié and Milan Vojnović. Coupon replication systems. In Proceedings of SIGMETRICS'05, volume 33, pages 2–13, Banff, Alberta, Canada, June 2005.
- [Ner81] Olle Nerman. On the convergence of supercritical general (C-M-J) branching processes. Z. Wahrscheinlichkeitstheorie verw. Gebiete, 57:365–395, 1981.
- [Nev86] J. Neveu. Erasing a branching tree. Advances in Applied Probability, (suppl.):101–108, 1986.
- [NQP08] R. Núñez-Queija and B. J. Prabhu. Scaling laws for file dissemination in P2P networks with random contacts. In *Proceedings of IWQoS*, 2008.
- [PWMC08] Nadim Parvez, Carey Williamson, Anirban Mahanti, and Niklas Carlsson. Analysis of bittorrent-like protocols for on-demand stored media streaming. In *Proceedings of SIGMETRICS'08*, pages 301–312, New York, NY, USA, 2008. ACM.
- [QS04] Dongyu Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proceedings of SIGCOMM'04*, pages 367–378, New York, NY, USA, August 2004. ACM.

- [Rob03] Philippe Robert. Stochastic Networks and Queues. Stochastic Modelling and Applied Probability Series. Springer-Verlag, New York, 2003. xvii+398 pp.
- [RS09] Philippe Robert and Florian Simatos. Occupancy schemes associated to Yule processes. *Advances in Applied Probability*, 41(2):600–622, June 2009.
- [RW87] L. C. G. Rogers and David Williams. Diffusions, Markov processes, and martingales. Vol. 2: Itô calculus. John Wiley & Sons Inc., New York, 1987.
- [SAV06] Riikka Susitaival, Samuli Aalto, and Jorma Virtamo. Analyzing the dynamics and resource usage of P2P file sharing by a spatio-temporal model. International Workshop on P2P for High Performance Comptutational Sciences (P2P-HPCS'06) in conjunction with ICCS, 2006.
- [Sen70] E. Seneta. On the supercritical branching process with immigration. Mathematical Biosciences, 7:9–14, 1970.
- [SRG08] Florian Simatos, Philippe Robert, and Fabrice Guillemin. A queueing system for modeling a file sharing principle. In *Proceedings of SIG-METRICS'08*, volume 36, pages 181–192, New York, NY, USA, 2008. ACM.
- [ST09] Florian Simatos and Danielle Tibi. Spatial homogenization in a stochastic network with mobility. *Annals of Applied Probability*, to appear, 2009.
- [Wil91] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.
- [YdV04] Xiangying Yang and Gustavo de Veciana. Service capacity of peer to peer networks. In INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies, volume 4, pages 2242–2252, 2004.

CHAPTER III

A Queueing System for Modeling a File-Sharing Principle

1.	Introduction	109		
2.	Model Description	110		
3.	Bins and Balls Problem	114		
4.	Deterministic Problem	116		
5.	Random Problem	122		
6.	Discussion	124		
References of Chapter III 1				

1. Introduction

This chapter analyzes the performance of a simple file sharing principle during a flash crowd scenario when a popular content becomes available on a peer-topeer network. It is supposed that a peer is willing to share a given file with a community of N peers, which are initially asleep. An asleep peer becomes active at some random time, i.e., it tries to download the file from a peer having the complete file. Once a peer has downloaded the file, it immediately becomes a server from which another peer can download the file. To simplify the model, we assume that the file is in one piece and not segmented into chunks; the time needed to download the file from one server is supposed to be random in order to take into account the diversity of upload capacities of peers. The goal of this chapter is to understand how the network builds up in this situation as peers join the system. In particular, we are interested in analyzing the growth of the number of available servers in the system. Note that there are eventually N + 1 servers since each peer eventually completes the file download.

In spite of its apparent simplicity, the analysis of the system is quite difficult because we have to cope with a network comprising a random number of servers: When peers complete their download, they become new servers so that the number of servers is continually increasing. It is assumed that an incoming peer chooses a server with the smallest number of queued peers. Other routing policies are considered at the end of the chapter.

The analysis performed in this chapter substantially differs from earlier studies appeared so far in the technical literature in the sense that we consider the transient formation of a network of peers. Yang and de Veciana [YdV06] considered a similar setting which they analyzed with results related to branching processes to describe the exponential growth of the number of servers. Our goal in this chapter is precisely to obtain more detailed asymptotics of this transient regime. Except the paper by Yang and de Veciana [YdV06], most of the papers published so far on the performance of peer-to-peer systems assume that peers join and leave the system and that a steady state regime exists. The problem is then to evaluate the impact of some parameters of the file sharing protocol on the equilibrium of the system. Different techniques can be used to perform such an analysis, for instance by using a Markovian chain to describe the state of the system, possibly by using approximation techniques when the state space related to the number of peers in the system is too large. See Ge *et al.* $[GFS^+03]$. A fluid flow analysis with an underlying Markovian structure is proposed in Clévenot and Nain [CN04] in order to model the Squirrel peer-to-peer caching system. In Qiu and Srikant [QS04], the authors directly use a fluid approximation to study the steady state of a peerto-peer network, subsequently complemented by diffusion variations around the steady state solutions. In Massoulié and Vojnović [MV05], the authors study the performance of a file sharing system via a stochastic coupon replication formulation, a coupon corresponding to a chunk of a file. The goal of this study is to understand the impact of the policy applied by users for choosing coupons on the performance of the system. The system is studied in equilibrium as in Qiu and Srikant [QS04].

The rest of this chapter is organized as follows: In Section 2, the system under consideration and some heuristics to study it are presented. It turns out that the its dynamics can be decomposed in two regimes. In the first one, there are almost no empty servers and we establish an analogy with a random bins and balls problem on the real line. By approximating the probability of selecting a bin by its mean value, we analyze in Section 3 the corresponding deterministic bins and balls problem. The analysis for the random bins and balls problem is much more complicated to analyze. The complete analysis is done in Chapter IV, corresponding to the paper Robert and Simatos [**RS09**], and the results useful for this chapter are summarized in Section 5. In Section 6, we support via simulation the different approximations and heuristics made in this chapter to analyze the file-sharing system, and we conclude this section by mentioning some possible extensions.

2. Model Description

Problem Formulation. We consider throughout this paper a system composed of N peers interested in downloading a given file. At the beginning, only one peer (the initial server) has the file and other peers are asleep. When becoming active, after an exponentially distributed duration of time with parameter ρ , a peer tries to download the file from the server that is the less loaded in terms of number of queued peers. In particular, the first peer becoming active downloads the file from the initial server. The time needed to download the file is assumed to be exponentially distributed with mean 1.

The hypothesis on the distribution on the duration of the time for a peer to become active is quite reasonable: this is a classical situation when a large number of *independent* users may access some network. The assumption on the duration of the time to download is not realistic in practice since this quantity is related to the size of the file requested whose distribution is more likely to be bounded by the maximal size of a chunk. As it will be seen, even within this simplified setting (in order to have a nice probabilistic description of the process), mathematical problems turn out to be quite intricate to solve. In this respect, our study could be seen as a first step in the analysis of flash crowd scenarios. It turns out that our current investigations in the general case seem to show that the exponential distribution does not have a critical impact on the qualitative behavior as long as the FIFO policy is used by servers. Mathematically, however, numerous technical points are not settled in this case.

We assume that peers requesting the file from the same server are served according to the FIFO discipline. Note that, because of the exponential distribution assumption, this case is equivalent to the Processor-Sharing discipline, i.e., when Npeers are present for a duration of time h, each of them receives the amount of work h/N. Just after completing the file download, a peer immediately becomes a server from which other peers can retrieve the file. The problem of "free riders", i.e., peers who do not become servers after service completion, is discussed briefly in Section 6. The conclusion is that this feature does not change significantly the qualitative properties of the system. The problem of servers who disconnect while they have downloads in progress will not be discussed in this paper.

It is worth noting that the model under consideration describes a "flash crowd" scenario. Indeed, a peer having a file accepts to share it with other peers and we are interested in the dynamics of the sharing process when a large population of peers tries to download the file. Moreover, since the durations for which these peers stay inactive are independent and identically distributed, the flow of arrivals of peers into the system is not stationary, but rather accumulates at the beginning and is then less and less intense. We are hence interested in the transient regime of the system. Contrary to the earlier studies [GFS+03, MV05, QS04], we are not interested in the steady state regime of the system, where peers continually join and leave the system.

It is intuitively clear that there should exist two different regimes for this system. Initially, it starts congested: many peers request the file, and only a few servers are available. Afterward, the situation is reversed: there are a large number of servers and only a few requests from the remaining inactive peers.

These two regimes clearly appear in Figure 1 depicting simulation results with $N = 10^6$ peers and $\rho = 5/6$. It shows that before time $T \approx 7$ time units (or equivalently mean download times), there are almost no empty servers, while after that time, more and more servers are empty until all peers have completed their download. But as long as the input rate is high, a new server immediately receives a customer. This is all the more true under the routing policy considered, since new peers entering the system choose an empty server if any.

A Non-Trivial Queueing Model. From the above description, the system can be represented by means of a queueing system with a random number of queues. Initially, the system is composed of a single server, and once a customer has completed its service, it becomes a new server. Since only a finite total number of customers is considered, there are eventually N + 1 servers.

When peer inter-arrival times and file download times are assumed to be exponentially distributed, a minimal Markovian representation of this queueing model requires the knowledge of the number of peers which are still asleep and the number of peers connected to each server. Since this Markov process is ultimately absorbing (all peers are servers at the end), the transient behavior of the system is of course



FIGURE 1. Results of a simulation: Fraction of idle servers against time for $N = 10^6$ and $\rho = 5/6$.

the main object of interest in the analysis. Even in very simple queueing systems, the transient behavior is delicate to analyze and much more difficult to describe than the stationary behavior. The classical M/M/1 queue is a good (and simple) example of such a situation when transient characteristics are not easy to express with simple closed form formulas. See Asmussen [Asm87] for example.

Given the multi-dimensional description (with unbounded dimension) of the Markov process, the system considered here is much more intricate and challenging. To analyze this system, a simpler mathematical model with bins and balls is used to investigate the duration of the first regime of this system. The specific point addressed in this paper is to describe the transient behavior when N becomes large.

Modeling the First Regime. Initially, the input rate is large and therefore a newly created server receives very quickly many requests from the numerous peers becoming active. The first regime described in the previous section and illustrated in Figure 1 is hence characterized by the fact that the duration times during which some servers are idle are negligible. In a second phase the number of empty servers begins to be significant before increasing very rapidly in the last phase. This phenomenon is discussed in Section 6. For the first regime, this leads us to describe the dynamics of the system as follows.

Let S_n be the time at which the *n*th server is created, with the convention that $S_0 = 0$ (the initial server has label 0). During the *n*th time interval (S_{n-1}, S_n) for $n \ge 1$, there are by definition exactly *n* servers. So if we neglect empty servers, $S_n - S_{n-1}$ is approximatively given by the minimum of *n* independent exponential random variables with parameter 1. The random variable S_n can thus be represented as $S_{n-1} + E_n^1/n$, where E_n^1 is an exponential random variable with parameter 1 independent of the past. In particular, during the first regime, the following approximation is accurate.

APPROXIMATION B. For $n \in \mathbb{N}$, as long as the system is still in the first regime, the instant of creation of the *n*th server is given by $S_n \approx T_n$, where

(1)
$$T_n = \sum_{k=1}^n \frac{E_k^1}{k},$$

and $(E_k^1, k \ge 1)$ being i.i.d. exponential random variables with unit mean.

The letter B used to designate this approximation stands for "branching": a Yule process is a special type of branching process where particles live for an exponential duration with mean 1, and split into two identical and independent particles upon death; see Athreya and Ney [AN72]. The previous discussion amounts precisely to say that as long as empty servers can be neglected, then the number of servers evolves like the population of a Yule process: each server, after a time exponentially distributed representing the service time, creates a new server. Equivalently, the old server dies and two new servers are created. The sequence (T_n) is called the sequence of *split times* of the Yule process.

Despite this approximation seems to be quite rough (a rigorous mathematical formulation of the approximation $S_n \approx T_n$ seems to be difficult to establish), Proposition 2.1 and the subsequent discussion below provide strong arguments to support its accuracy. Approximation B only tells about the process of creation of servers when empty servers can be neglected: it does not tell anything about the duration of the first regime, i.e., the time until which this approximation indeed holds. We rely on a heuristic to determine this duration; the heuristic we choose in Definition 1 is discussed in light of our analytical results of Sections 4 and 5 in Section 6.

Figure 1 suggests that the last time there is no empty server closely coincides with the end of the first regime that we want to estimate. This time is unfortunately not a stopping time and turns out to be much more difficult to study; moreover Approximation B, which is the fundamental assumption underlying our analysis, seems to fail long before this time (see Section 6). Rather, we choose our heuristic according to the following definition; different choices are discussed in Section 6.

DEFINITION 1. The duration of the first regime is defined as S_{ν} , where ν is the first index $n \ge 1$ so that one or no peer arrive between S_{n-1} and S_n .

According to this definition, the first regime lasts as long as between the creation of two successive servers, at least two peers arrive in the system. The intuition behind this heuristic is that, because of the policy for the choice of servers, if many peers arrive in any interval, then the least loaded servers will receive requests from arriving peers. Thus, as long as many peers arrive, it is quite rare for a server to remain empty.

The phase transition should occur when the number of arrivals between the creation of two successive servers is not sufficient to give work to empty servers which are created. In particular, if no peers arrive in some interval, then there will be at least two empty servers at the beginning of the next time interval. So the first time when only a few peers arrive in some interval should be a good indication on the current state of the system. A probably more natural heuristic would have been to consider the first interval in which no peer arrives. Nevertheless, an argument in favor of the former heuristic is that it enjoys the following nice property.

PROPOSITION 2.1. For $n < \nu$, at most two servers are simultaneously empty in the *n*th interval (S_{n-1}, S_n) .

PROOF. The proof is by induction. For n = 1, the property is trivial, since there is only one server in the first interval. Consider now $1 < n < \nu$, and suppose that the property holds for n - 1. Since at least two peers arrive in the (n - 1)th interval, and since these peers are necessarily routed to empty servers, if any, there is no empty server just before S_{n-1} . Therefore, just after S_{n-1} , there are at most two empty servers, and so the property holds as long as $n < \nu$.

We are now able to justify Approximation B for $n < \nu$. The argument is twofold, depending on whether n is small or large. In the former case, the *time* during which servers are empty is neglegible, while in the latter the *fraction* of empty servers is negligible.

Indeed, for any $n < \nu$ the number of non idle servers is between n-2 and n as shown by Proposition 2.1. Hence the fraction of empty servers is negligible for n large, and $S_n - S_{n-1}$ should thus be close in distribution to an exponentially distributed random variable with parameter n.

Now for n bounded, there may be at any time 2 empty servers, and this may represent some fraction of the total number of servers. However the time during which there are empty servers is small. For n fixed, it is easy to see that the mean number of peers that arrive in the nth interval is indeed of order of N. Hence right after the nth server is created, the next peer arrives very rapidly, after a time of order of 1/N, which is arbitrarily large when N is large. Hence servers cannot remain empty for a long time.

From now on, the identification of S_n and T_n , where the sequence (T_n) is defined by Equation (1), is assumed to hold. Results on T_n can be assumed to hold for S_n when $n < \nu$.

3. Bins and Balls Problem

Denote by $(E_i^{\rho}, 1 \leq i \leq N)$ a sequence of i.i.d. random variables, exponentially distributed with parameter ρ . For $i \leq N$, E_i^{ρ} is the time at which the *i*th peer becomes active.

We introduce the following bins and balls model on the real line: The interval (T_{n-1}, T_n) is the *n*th bin and the variables $(E_i^{\rho}, 1 \le i \le N)$ are the locations of N balls thrown on the real line. The set $\{T_{n-1} \le E_i^{\rho} \le T_n\}$ is simply the event that the *i*th ball falls into the *n*th bin. Conditionally on the sizes of the bins, i.e., on $\mathcal{T} = (T_n)$, the probability of such an event (which does not depend on *i*) is

(2)
$$P_n = \mathbb{P}\left(T_{n-1} < E_i^{\rho} < T_n \mid \mathcal{T}\right) = e^{-\rho T_{n-1}} \left(1 - e^{-\rho E_n^1/n}\right),$$

where the random variables E_n^1 , $n \ge 1$, are independent and exponentially distributed with mean unity.

With the above formulation, we have then to deal with the following bins and balls model:

- (1) A random probability distribution $\mathcal{P} = (P_n)$ is given (bins with random sizes).
- (2) N balls are thrown independently according to the probability distribution \mathcal{P} .

It is worth noting that the above bins and balls model has an infinite number of bins. In addition, although bins and balls problems have been widely studied in the literature, our model presents a remarkable feature: For $i \ge 1$, a ball falls into bin *i* with probability P_i which is a random variable, but conditionally on the sequence (P_n) , this is a classical bins and balls problem. Mathematical results for bins and balls models with random distributions are quite rare. See Kingman [Kin78] and Gnedin *et al.* [GHP07] and the references therein where some related models have been investigated.

The random model under consideration will give us some information on the behavior of our system. The following proposition establishes a simple but important characterization for the asymptotic behavior of (P_n) .

PROPOSITION 3.1. Let $(E_i^1, i \ge 1)$ be independent exponential random variables with parameter 1. Then, for $n \in \mathbb{N}$

(3)
$$T_n = \sum_{k=1}^n \frac{E_k^1}{k} \stackrel{dist.}{=} \max_{1 \le k \le n} E_k^1;$$

and the sequence $(T_n - \log n)$ converges almost surely to a finite random variable T_{∞} whose distribution is given by $\mathbb{P}(T_{\infty} \leq x) = \exp(-\exp(-x))$ for $x \in \mathbb{R}$.

The conditional probability P_n of throwing a ball into the nth bin can be written as

$$P_n = \frac{\rho}{n^{\rho+1}} X_{n-1} Z_n,$$

where

$$Z_n = \frac{n}{\rho} \left(1 - e^{-\rho E_n^1/n} \right) \text{ and } X_{n-1} = n^{\rho} e^{-\rho T_{n-1}}$$

are independent random variables. As n goes to infinity, X_n (resp. Z_n) converges in distribution to X_{∞} (resp. Z_{∞}). The convergence of (X_n) to X_{∞} holds almost surely and in L_q , for any $q \ge 1$.

The limiting variable Z_{∞} has an exponential distribution with parameter 1 and X_{∞} has a Weibull distribution with parameter $1/\rho$,

(5)
$$\mathbb{P}(X_{\infty} \ge x) = e^{-x^{1/\rho}}, \ x \ge 0$$

PROOF. Let $E_{(1)} \leq E_{(2)} \leq \cdots \leq E_{(n)}$ be the variables $(E_k^1, 1 \leq k \leq n)$ in increasing order. In particular $E_{(n)} = \max_{1 \leq k \leq n} E_k^1$. With the convention $E_{(0)}=0$, due to standard properties of the exponential distribution, the variables $E_{(i+1)}-E_{(i)}$, $i = 0, \ldots, n-1$ are independent and the variable $E_{(i+1)} - E_{(i)}$ is the minimum of n - i exponential variables with parameter 1, i.e., has the same distribution as $E_{n-i}^1/(n-i)$. The distribution identity (3) then follows.

Since $Z_n \stackrel{\text{dist.}}{=} n/\rho(1 - \exp(-\rho E_1/n))$, it converges in distribution to an exponential distribution with parameter one.

Define

$$M_n = \sum_{k=1}^n \frac{E_k^1 - 1}{k} = T_n - H_n,$$

where (H_n) is the sequence of harmonic numbers, $H_n = 1 + 1/2 + \cdots + 1/n$. The sequence (M_n) is clearly a martingale, it is bounded in L_2 since

$$\mathbb{E}M_n^2 = \sum_{k=1}^n \frac{\mathbb{E}\left(E_k^1 - 1\right)^2}{k^2} = \sum_{k=1}^\infty \frac{1}{k^2} < +\infty.$$

It therefore converges almost surely. See Williams [Wil91] for example. The almost sure convergence of $(T_n - \log n) = (M_n + H_n - \log n)$ is thus proved. Identity (3) gives that, for $x \ge 0$,

$$\mathbb{P}(T_n - \log n \le x) = (1 - e^{-x - \log n})^n \sim e^{-e^{-x}},$$

as n goes to infinity.

Since

$$X_m = e^{-\rho M_n} e^{\rho(\log(n+1) - H_n)}$$

one gets the almost sure convergence of (X_n) . It is easy to check that, for $q \ge 0$,

(6)
$$\mathbb{E}\left(X_{n}^{q}\right) = (n+1)^{q\rho} \prod_{i=1}^{n} \frac{1}{1+q\rho/i} = (n+1)^{q\rho} \frac{\Gamma(n)}{\Gamma(n+q\rho)} \Gamma(q\rho) \sim \Gamma(q\rho),$$

when $n \to \infty$, where Γ is the usual Gamma function, and where the last equivalence easily comes from Stirling's Formula. In particular, for any $q \ge 0$, the *q*th moment of X_n is therefore bounded with respect to *n*. One deduces the convergence in L_q of the sequence (X_n) . Since $X_n = \exp(-\rho(T_n - \log(n+1)))$, one has the equality in distribution $X_{\infty} = \exp(-T_{\infty})$ which gives the law of X_{∞} .

It is important to note that the probability distribution $\mathcal{P} = (P_n)$ is a random element in the set of probability distributions on \mathbb{N} . The decay of this distribution follows a power law with parameter $\rho + 1$, because according to the previous proposition, $n^{\rho+1}P_n$ converges in distribution to $\rho X_{\infty} Z_{\infty}$. Using the asymptotic behavior derived in (6) with q = 1, it is easy to see that the average probability for a ball to fall into the *n*th bin satisfies the following relation

(7)
$$\mathbb{E}(P_n) \sim \frac{\rho \Gamma(\rho)}{n^{\rho+1}}.$$

This equivalence suggests the introduction of a deterministic version of the bins and balls problem considered.

4. Deterministic Problem

Description. Denote by $\mathcal{Q} = (q_n)$ a probability distribution on \mathbb{N} such that

(8)
$$\lim_{n \to +\infty} n^{\delta} q_n = \alpha,$$

for some $\alpha > 0$ and $\delta > 1$. For each n, q_n can be seen as the probability for a ball to fall in the *n*th bin. When $\delta = \rho + 1$ and $\alpha = \rho \Gamma(\rho)$, the sequence (q_n) has the same asymptotic behavior as $\mathbb{E}(P_n)$ given by Equation (7). Hence, this model may be considered as the deterministic equivalent of the bins and balls problem defined in the previous section. For the sake of clarity, the problem with the probability distribution \mathcal{P} (resp. \mathcal{Q}) will be referred to as the random (resp. deterministic) problem.

The deterministic problem amounts to throwing N exponential variables with parameter ρ on the half-real line, where this line has been divided into deterministic intervals (t_{n-1}, t_n) with $t_n = \mathbb{E}T_n$. The main quantity of interest in the following is the asymptotic behavior with respect to N of the index of the first bin that does not receive any ball. DEFINITION 2. Let us denote by $\eta_i^R(N)$ (resp. $\eta_i^D(N)$) the number of balls in the *i*th bin when N balls have been thrown in the random (resp. deterministic) bins and balls problem, and define

(9)
$$\nu^{R}(N) = \inf\{i \ge 1 : \eta^{R}_{i}(N) = 0\}, \\ \nu^{D}(N) = \inf\{i \ge 1 : \eta^{D}_{i}(N) = 0\}.$$

In view of Definition 1, to investigate the duration of the first regime of the system, the asymptotic behavior of the sequences $(\nu^R(N))$ and $(\nu^D(N))$ is analyzed. Since we consider that the first regime lasts until one or no peers arrive between the creation of two successive servers, we should have to consider $\nu'(N) = \inf\{i \ge 1 : \eta_i(N) \le 1\}$ to be rigorous. In fact, the mathematical analysis of the index of the first empty bin can easily be extended to the first bin that receives less than k balls, see the remark following Theorem IV.4.1. For the sake of simplicity, we therefore only treat the case k = 0. Neither the orders of magnitude nor the asymptotic behaviors established in the following are affected by the value of k, and in particular if we consider 1 instead of 0.

To conclude this section, let us give a rough approximation of the correct order of magnitude for $\nu^R(N)$ and $\nu^D(N)$ as N gets large. Rigorous mathematical analysis is carried out in Section 4, while Section 6 compares the insights provided by the two models.

For $i \geq 1$, $\mathbb{E}(\eta_i^D(N)) = Nq_i \sim \alpha N/i^{\rho+1}$. Hence, in the deterministic model, a finite number of balls will fall in the *i*th bin as soon as *i* is of the order of $N^{1/(\rho+1)}$ as *N* becomes large. Hence we expect that in the deterministic model, $\nu^D(N)/k(N)$ converges in distribution for $k(N) = N^{1/(\rho+1)}$. Theorem 4.1 below shows that the location of the first empty bin is in fact slightly smaller than $N^{1/(\rho+1)}$, i.e., of the order of $(N/\log N)^{1/(\rho+1)}$. Nevertheless this heuristic approach gives the correct exponent in *N*.

Although $\mathbb{E}(\eta_i^R(N))$ has the same asymptotic behavior, the corresponding heuristic approach in the case of the random model is more subtle. Indeed, we have

$$\mathbb{E}(\eta_i^R(N)) = N\mathbb{E}(P_i) \sim N\rho\Gamma(\rho)/i^{\rho+1},$$

so the number of balls falling in the *i*th bin should be of the order $Ni^{-\rho-1}$. However, in the random model, the *i*th interval is with random length E_i^1/i . So from T_{i-1} , the next point T_i is at a distance E_i^1/i and the first ball is at a distance corresponding to the minimum of $Ni^{-\rho-1}$ i.i.d. exponential random variables with parameter 1. Thus, with this approximation, the *i*th interval is empty with probability

$$\mathbb{P}\left(\frac{E_i^1}{i} \le \frac{i^{\rho+1}}{N} E_0^1\right) = \frac{1}{1 + N/i^{\rho+2}}.$$

When $N \to \infty$, this probability is non negligible as soon as *i* is of order $N^{1/(\rho+2)}$, which is significantly below what we found in the deterministic case. Theorem 5.1 below shows that this is indeed the correct answer. The order of magnitude is one order smaller, compared to the deterministic case, because of the variability of the intervals size: to some extent, a very small interval is generated, so that no balls fall in it, while in the deterministic case, some balls would have.

Asymptotic Analysis. Csáki and Földes [CF76] gives the asymptotic behavior of the distribution of ν^D when N is large. A more complete description of the locations of the first empty bins (and not only for the first one) can however be achieved. For this purpose, the variable W_N^k is defined as the number of empty bins whose index is less than k when N balls have been thrown. This random variable is formally defined as

(10)
$$W_N^k = \sum_{i=1}^k I_{N,i}, \text{ with } I_{N,i} = \mathbb{1}_{\{\eta_i^D(N)=0\}}.$$

The distribution of W_N^k is analyzed when k is dependent on N. First, some estimates for the mean value and the variance of W_N^k are required.

PROPOSITION 4.1. Assume that the sequence (q_i) is non-increasing. For x > 0, if

(11)
$$\kappa_x(N) = \left\lfloor \left(\alpha \delta \frac{N}{\log N} \right)^{1/\delta} \left[1 + \frac{1+\delta}{\delta} \frac{\log \log N}{\log N} + \frac{\log x}{\log N} \right] \right\rfloor$$

where $\lfloor y \rfloor$ is the integral part of y > 0, then

(12)
$$\lim_{N \to +\infty} \mathbb{E}\left(W_N^{\kappa_x(N)}\right) = (\alpha \delta)^{1/\delta} x.$$

PROOF. For $k, N \in \mathbb{N}$

(13)
$$\mathbb{E}\left(W_N^k\right) = \sum_{i=1}^{\kappa} (1-q_i)^N.$$

For $0 \le x \le 1$,

$$0 \le e^{-Nx} - (1-x)^N \le x_N (1-x_N)^{N-1},$$

where x_N is the unique solution to the equation $\exp(-Nx) = (1-x)^{N-1}$, since the function $x \to e^{-Nx} - (1-x)^N$ has a maximum at point x_N . It is easily seen that $Nx_N \leq 2$ (in fact $Nx_N \to 2$ as $N \to +\infty$), so that for $N \geq 1$

(14)
$$\sup_{0 \le x \le 1} \left| e^{-Nx} - (1-x)^N \right| \le \frac{2}{N}.$$

With this relation, we obtain

$$\left| \mathbb{E} \left(W_N^k \right) - \sum_{i=1}^k e^{-Nq_i} \right| \le \frac{2k}{N},$$

so that for $k = \kappa_x(N)$ and large N, $(1 - q_i)^N$ can be replaced with $\exp(-Nq_i)$ in the expression of $\mathbb{E}(W_N^k)$.

For the sake of simplicity, we assume that $q_i = \alpha/i^{\delta}$, for $i \ge 1$. The general case of a non-increasing sequence (q_i) follows along the same lines since the crucial relation below holds true with a convenient function q. One defines $q(x) = \alpha \min(x^{-\delta}, 1)$ for $x \ge 0$.

$$\int_0^k e^{-Nq(u)} \, du \le \sum_{i=1}^k e^{-Nq_i} \le \int_1^{k+1} e^{-Nq(u)} \, du.$$

The difference between these two integrals is bounded by $2 \exp(-\alpha N/k^{\delta})$. Now take k = k(N) with k(N) with the same order of magnitude as $(N/\log N)^{1/\delta}$, say, $k(N) \sim A(N/\log N)^{1/\delta}$ for some A > 0. We have

$$\mathbb{E}\left(W_N^{k(N)}\right) = \int_1^{k(N)} e^{-Nq(u)} \, du + o(1).$$

The right hand side of the above equation is given by

(15)
$$\int_{1}^{k(N)} e^{-\alpha N u^{-\delta}} du = \frac{(\alpha N)^{1/\delta}}{\delta} \int_{\alpha N k(N)^{-\delta}}^{\alpha N} e^{-u} u^{-(\delta+1)/\delta} du.$$

Now let $H(N) = \alpha N k(N)^{-\delta}$ and consider

$$e^{H(N)}H(N)^{(1+\delta)/\delta} \int_{H(N)}^{\alpha N} e^{-u} u^{-(\delta+1)/\delta} du$$

= $\int_{H(N)}^{\alpha N} e^{-(u-H(N))} \left(\frac{H(N)}{u}\right)^{-(\delta+1)/\delta} du$
= $\int_{1}^{\alpha N/H(N)} H(N) e^{-H(N)(u-1)} \frac{1}{u^{(\delta+1)/\delta}} du$
 $\sim \int_{0}^{+\infty} H(N) e^{-H(N)u} \frac{1}{(1+u)^{(\delta+1)/\delta}} du \sim 1$

since $N/H(N) \to +\infty$ and $H(N) \to +\infty$ as $N \to +\infty$. Therefore, an equivalent expression of the integral in the right hand side of Equation (15) has been obtained. A careful analysis shows that

$$\liminf_{N \to +\infty} \left(N^{1/\delta} e^{-H(N)} H(N)^{-(1+\delta)/\delta} \right) > 0$$

for $k(N) = \kappa_x(N)$, therefore the leading term in $\mathbb{E}(W_N^{\kappa_x(N)})$ is indeed the integral given by (15). Gathering these results, we obtain

(16)
$$\mathbb{E}\left(W_N^{\kappa_x(N)}\right) = \frac{(\alpha N)^{1/\delta}}{\delta} e^{-H(N)} H(N)^{-(1+\delta)/\delta} + o(1)$$
$$\sim \frac{1}{\alpha\delta} \frac{k(N)^{1+\delta}}{N} \exp\left(-\alpha N \kappa_x(N)^{-\delta}\right)$$

from which Relation (12) is obtained.

The following proposition shows the equivalence of the variance and the mean value of $W_N^{\kappa_x(N)}$ under a convenient scaling. This result is crucial to prove the limit theorems of this section.

PROPOSITION 4.2. Assume that the sequence (q_i) is non-increasing. For x > 0, let κ_x be defined by Equation (11), then

(17)
$$\lim_{N \to +\infty} \operatorname{Var}\left(W_N^{\kappa_x(N)}\right) / \mathbb{E}\left(W_N^{\kappa_x(N)}\right) = 1$$

PROOF. For $k \ge 1$, by using Equation (13) (which does not depend on α).

$$(\mathbb{E}[W_N^k])^2 = \sum_{1 \le i,j \le k} (1 - q_i - q_j + q_i q_j)^N,$$

and

$$\mathbb{E}[(W_N^k)^2] = \mathbb{E}[W_N^k] + \sum_{1 \le i \ne j \le k} (1 - q_i - q_j)^N,$$

so that, to prove the equivalence of $\mathbb{V}ar(W_N^{\kappa_x(N)})$ and $\mathbb{E}(W_N^{\kappa_x(N)})$, it is sufficient to show that the quantities

$$\sum_{1 \le i,j \le \kappa_x(N)} \left[(1 - q_i - q_j + q_i q_j)^N - (1 - q_i - q_j)^N \right] \text{ and } \sum_{i=1}^{\kappa_x(N)} (1 - 2q_i)^N$$

are negligible with respect to $\mathbb{E}(W_N^{\kappa_x(N)})$. This amounts to show that these quantities are o(1) by Proposition 4.2. The second term is the expected number of empty bins for the distribution (\tilde{q}_i) such that $\tilde{q}_i \sim 2\alpha/i^{\delta}$. Estimate (16) shows that

$$\sum_{i=1}^{\kappa_x(N)} (1-2q_i)^N \sim \frac{1}{2\alpha\delta} \frac{\kappa_x(N)^{1+\delta}}{N} \exp\left(-2\alpha N\kappa_x(N)^{-\delta}\right) = o\left(\mathbb{E}W_N^{\kappa_x(N)}\right).$$

By using the fact that for $a \ge b \ge 0$, $a^N - b^N \le N(a - b)a^{N-1}$, the second term satisfies

(18)
$$\sum_{1 \le i,j \le k} \left[(1 - q_i - q_j + q_i q_j)^N - (1 - q_i - q_j)^N \right]$$
$$\le N \sum_{1 \le i,j \le k} q_i q_j (1 - q_i - q_j + q_i q_j)^{N-1} = \frac{1}{N} \left(\sum_{i=1}^k N q_i (1 - q_i)^{N-1} \right)^2.$$

By using a similar method as in the proof of Proposition 4.1, we obtain the equivalence

$$\sum_{i=1}^{k(N)} Nq_i (1-q_i)^{N-1} \sim \int_1^{k(N)} Nq(u) e^{-Nq(u)} du$$
$$\sim (\alpha\delta)^{1/\delta} \alpha x \frac{N}{\kappa_x(N)^\delta} = (\alpha\delta)^{1/\delta} x \log N$$

This equivalence together with Equation (18) complete the proof of the proposition.

THEOREM 4.1. Let (q_n) be a non-increasing sequence satisfying Relation (8). For x > 0 and $N \in \mathbb{N}$, set

$$\kappa_x(N) = \left\lfloor \left(\alpha \delta \frac{N}{\log N} \right)^{1/\delta} \left(1 + \frac{1+\delta}{\delta} \frac{\log \log N}{\log N} + \frac{\log x}{\log N} \right) \right\rfloor$$

When N goes to infinity, the variable $W_N^{\kappa_x(N)}$ converges in distribution to a Poisson random variable with parameter $(\alpha \delta)^{1/\delta} x$.

The index $\nu^D(N)$ of the first empty bin defined by Equation (9) is such that the variable

(19)
$$\frac{(\log N)^{(1+\delta)/\delta}}{(\alpha\delta N)^{1/\delta}}\nu^D(N) - \log N - \frac{1+\delta}{\delta}\log\log N$$

converges in distribution to a random variable Y defined by

$$\mathbb{P}(Y \ge x) = \exp\left(-(\alpha \delta)^{1/\delta} e^x\right), \quad x \in \mathbb{R}.$$

PROOF. Chen-Stein's method is the basic tool in the proof of the theorem. See Barbour *et al.* [**BHJ92**] for a detailed presentation of this powerful method. Let N and k be in \mathbb{N} and $1 \leq i_0 \leq k$. The variable W_N^k conditioned on the event $\{I_{N,i_0} = 1\}$ has the same distribution as the number of empty bins when the balls in the i_0 th bin are thrown again until the i_0 th bin is empty. It follows that the number of balls in any other bin is larger than in the case when they are assigned at first draw. One deduces that for $i \neq i_0$,

$$\mathbb{P}(I_{N,i} = 1 \mid I_{N,i_0} = 1) \leq \mathbb{P}(I_{N,i} = 1).$$

The variables $(I_{N,i}, 1 \leq i \leq k)$ are therefore negatively correlated, see Barbour *et al.* [BHJ92]. Then, by [BHJ92, Corollary 2.C.2], the following relation holds,

$$\sum_{p\geq 0} \left| \mathbb{P}(W_n^k = p) - \frac{\mathbb{E}\left(W_N^k\right)^p}{p!} e^{-\mathbb{E}\left(W_N^k\right)} \right| \le 1 - \mathbb{V}\mathrm{ar}\left(W_N^k\right) / \mathbb{E}\left(W_N^k\right).$$

By taking $k = \kappa_x(N)$ and by using Propositions 4.1 and 4.2, we obtain the convergence in distribution of $W_N^{\kappa_x(N)}$ to a Poisson distribution with parameter $(\alpha\delta)^{1/\delta}x$. The last part of the theorem is a simple consequence of the identity $\mathbb{P}(W_N^k = 0) = \mathbb{P}(\nu^D(N) > k)$.

The convergence in distribution of $\nu^D(N)$ has been proved previously by Csáki and Földes [CF76] with a different method. Our result gives a more accurate description of the location of empty bins (and not only the first one) near the index $\kappa_x(N)$.

The following corollary is a straightforward application of the detailed asymptotics obtained in the above theorem.

COROLLARY 4.1 (Cutoff phenomenon). Under the assumption of Theorem 4.1, if

$$k(N) = \left(N/\log N\right)^{1/\delta}$$

then, as N goes to infinity, the following convergence in distribution holds: For $\beta > 0$,

$$W_N^{\beta k(N)} \longrightarrow \begin{cases} +\infty & \text{if } \beta > (\alpha \delta)^{1/\delta}, \\ 0 & \text{if } \beta < (\alpha \delta)^{1/\delta}. \end{cases}$$

So far, only indexes of empty bins have been considered. The result below shows that the first empty bin is located on the time axis at a time of the order of $\log N$. It will be discussed in Section 6 why this suggests that the time the system begins to serve quickly the incoming peers should be of the same order.

COROLLARY 4.2 (First Empty Bin). Let

(20)
$$T^{D}(N) = T_{\nu^{D}(N)} = \sum_{k=1}^{\nu^{D}(N)} \frac{E_{k}^{1}}{k}$$

Under the assumptions of Theorem 4.1, the quantity

$$\delta T^D(N) - \log N + \log \log N - \log(\alpha \delta)$$

converges in distribution to δT_{∞} , where T_{∞} is the random variable defined in Proposition 3.1.

PROOF. We have

$$\delta T^{D}(N) - \log N + \log \log N - \log(\alpha \delta)$$

= $\delta \left(\sum_{k=1}^{\nu^{D}(N)} \frac{E_{k}^{1}}{k} - \log(\nu^{D}(N)) \right) + \log \left(\frac{(\nu^{D}(N))^{\delta} \log N}{\alpha \delta N} \right).$

Since $(T_N - \log N)$ converges almost surely to T_∞ and $(\nu^D(N))^{\delta} \log N/N$ converges in distribution to 1 by Theorem4.1, Skorokhod representation theorem shows that one can assume that the two-dimensional sequence $(T_N - \log N, (\nu^D(N))^{\delta} \log N/N)$ converges almost surely to $(T_\infty, 1)$. In view of the right-hand side of the last display, one gets that on this probability space the sequence $(\delta T^D(N) - \log N + \log \log N - \log(\alpha \delta))$ converges almost surely to δT_∞ , hence the convergence holds in law as well, independently of the probability space.

5. Random Problem

For the random model, the probability P_n of selecting the *n*th bin is given by Equation (4) of Proposition 3.1. As *n* goes to infinity, $X_n \sim X_\infty$ almost surely and in distribution, Z_n is asymptotically an exponentially distributed random variable with parameter 1. The sequence (P_n) can be approximated by

$$\left(\frac{\rho}{n^{\rho+1}}X_{\infty}E_n^1\right)$$

where (E_n^1) are i.i.d. exponential variables with unit means.

In spite of the fact that the decay of P_n follows a power law, the random factor plays an important role. This factor is composed of two variables, one (namely X_{∞}) is fixed once for all and the other (namely Z_n) changes for every bin. The fact that Z_n , related to the "width" of the *n*th bin, can be arbitrarily small with a positive probability suggests that the index ν^R of the first empty bin should be smaller than the corresponding quantity for the deterministic case. This is indeed true but the situation in this case is much more complex to analyze. The complete analysis of the random case is given in the next chapter, and only the results relevant to our problem are summarized here. It must be noticed that a similar problem where X_{∞} and the sequence $(E_n^1, n \ge 1)$ are independent is fairly easy to solve. However here, these random variables are dependent, and this dependency requires quite technical probabilistic tools.

The asymptotic result on ν^R of Corollary 5.1 is a reformulation of Corollary IV.4.1 with the notations of this section. The results of Chapter IV deal with point processes which give comprehensive information on the locations of the first empty bins; however in order to stick with the results of the previous section, we rephrase these results in terms of the random variable that counts the number of empty bins. Hence the asymptotic behavior of the random variable W_N^k defined by

$$W_N^k = \sum_{i=1}^k I_{N,i}$$
 with $I_{N,i} = \mathbb{1}_{\{\eta_i^R(N)=0\}}$.

is investigated. Although in the deterministic case, Chen-Stein's method makes it possible to reduce the analysis of W_N^k to its first and second moments, this is no longer the case for the random problem. Indeed, because of the variability of the bins sizes, the random variables $(I_{N,i}, 1 \leq i \leq k)$ are no longer negatively correlated. Moreover, the ratio of the expected value to the variance of $W_N^{k(N)}$ does not converge to 1 for a convenient sequence (k(N)) as in the deterministic case (Proposition 4.2), which suggests that if a limit in distribution exists, it cannot be Poisson.

As was pointed out in Hwang and Janson [HJ08], the sequence $(NP_i, 1 \le i \le k)$ plays a central role in the limiting behavior of (W_N^k) . The intuitive explanation, developed in the next chapter, is that the first bins with index *i* such that NP_i is of order of 1 are actually the first empty bins. The following technical proposition gives a result on the asymptotic behavior of this sequence. It is important since it introduces the scale $N^{1/(\rho+2)}$ which turns out to be the correct scaling for the variable $\nu^R(N)$; moreover, it highlights the sequence $(NP_i, 1 \le i \le xN^{1/(\rho+2)})$ which is important. In the remainder of this section, denote

$$\kappa(N) = N^{1/(\rho+2)}.$$

PROPOSITION 5.1. Let x > 0. When N goes to infinity, the random sequence $(NP_i, 1 \le i \le x\kappa(N))$ converges in distribution to a mixed Poisson process with a random intensity distributed like $x^{\rho+2}(X_{\infty}\rho(\rho+2))^{-1}$.

PROOF. This is a consequence of Theorem IV.3.1 and of an approximation technique used in the proof of Proposition IV.3.3. The proof remains very concise in order to avoid unnecessary repetitions of technical arguments which fit better in the next chapter. Let $g : \mathbb{R}_+ \to \mathbb{R}_+$ be a continuous function with a compact support: the theorem will be implied by the convergence of Laplace transforms:

$$\lim_{N \to +\infty} \mathbb{E}\left(e^{-\sum_{i=1}^{x \kappa(N)} g(NP_i)}\right) = \mathbb{E}\left(e^{-x^{\rho+2}X_{\infty}^{-1}/(\rho(\rho+2))}\right)$$

By defining the function $f(u, v) = \mathbb{1}_{\{0 \le u \le x\}} g(v)$, we see that this convergence is the same as the convergence (IV.7) but with a function f which is not continuous. Approximating f with continuous functions similarly as in the proof of Proposition IV.3.3 then gives the result.

This result together with standard poissonization techniques (and again the approximation technique of the proof of Proposition IV.3.3) make it possible to prove the following theorem, which is the main result of this section; see the proof of Theorem IV.4.1 for details on poissonization techniques.

THEOREM 5.1. Let x > 0. When N goes to infinity, $W_N^{x\kappa(N)}$ converges in distribution to a mixed Poisson random variable with parameter $x^{\rho+2} (X_{\infty}\rho(\rho+2))^{-1}$.

COROLLARY 5.1. When N goes to infinity, the random variable $\nu^R(N)/\kappa(N)$ converges in distribution to a random variable Y such that

$$\mathbb{P}(Y \ge x) = \mathbb{E}\left(e^{-x^{\rho+2}X_{\infty}^{-1}/(\rho(\rho+2))}\right).$$

Finally, if $T^R(N) \stackrel{\text{def.}}{=} T_{\nu^R(N)}$ then, for the convergence in distribution,

(21)
$$\lim_{N \to +\infty} \frac{T^R(N)}{\log(N)} = \frac{1}{\rho + 2}$$

PROOF. The asymptotic behavior of $\nu^R(N)$ follows from Theorem 5.1 and the equality

$$\mathbb{P}(\nu^R(N) > n) = \mathbb{P}(W_N^n = 0)$$

The asymptotic behavior of the sequence $(T^R(N))$ follows similarly as in the proof of Corollary 4.2.

The fact that the parameter of the limiting Poisson law is random has important effects, especially concerning the expectation. Indeed, it stems from Equation (5) and Proposition 3.1 that $\lim_{N\to\infty} \mathbb{E}(W_N^{\kappa\kappa(N)})$ is proportional to $\mathbb{E}(X_\infty^{-1})$ and, from Proposition 3.1, $\mathbb{E}(X_\infty^{-1}) < +\infty$ if and only if $\rho < 1$. Note in particular that the value $\rho = 1$ plays a special role for our system. For $\rho > 1$, the mean value of $W_N^{\kappa\kappa(N)}$ diverges because it happens that a finite number of intervals (actually, the $\lfloor \rho \rfloor$ first intervals) capture most of the balls. This

For $\rho > 1$, the mean value of $W_N^{x_{N(N)}}$ diverges because it happens that a finite number of intervals (actually, the $\lfloor \rho \rfloor$ first intervals) capture most of the balls. This event happens with an increasingly small probability, so that in the limit as N goes to infinity, it does not have any impact on our system. However, for a fixed N, this event happens with a fixed probability as well. For instance, we commonly observed on various simulations for $\rho = 2$ and N = 10000 that more than 95% of the peers go to the first server, which is clearly an undesirable behavior of the system.

6. Discussion

In this section, a set of simulations of the file-sharing principle is presented to test the different approximations made in this chapter in terms of bins and balls. These simulations make it possible to confirm three main points:

- The population of servers initially resembles a Yule process, i.e., the branching Approximation B holds for some time.
- Approximation B holds until the time T^D predicted in Corollary 4.2 by the deterministic bins and balls problem.
- Approximation B does not hold after T^D .

Since the deterministic model predicts a first empty bin later than in the random model, the second point entails that the deterministic model is more accurate than the random one, and we explain in the following why. Since the branching Approximation B is the cornerstone of our analysis, the last point suggests that a different technical approach must be used to study the system after T^D .

We moreover use these simulations to discuss heuristics different than the one given by Definition 1 and considered so far, as well as a different server selection policy, which consists in letting an incoming peer choose the server at random. A surprising observation made from these simulations is that the branching Approximation B does not hold until empty servers constitute a significant proportion of the servers: it breaks down even though non-empty servers are still prevailing.

Throughout this section, we discuss the relevance of several random variables. The goal is to assess the accuracy of the procedure consisting of estimating the length of the first regime by using the random variable ν specified in Definition 1. For this purpose, three different times H_1 , H_2 and H_3 are discussed based on simulations:



FIGURE 2. $\log(\mathbb{E}(\nu_1))$ (solid) and $\mathbb{E}(H_1)$ (dashed) against $\log N$.

Heuristic (1): H_1 is the first time when two servers are created and less than 2 peers have arrived.

Heuristic (2): H_2 is the last time when there is an empty server.

Heuristic (3): H_3 is the first time when a server becomes empty, i.e., when a peer leaves a server where it was alone.

Note that the difference between H_1 and T^R is precisely the branching Approximation B: in the definition of H_1 , empty servers are not neglected. We also discuss theoretical results linked with a fourth heuristic:

Heuristic (4): H_4 is the first time when the input rate is smaller than the output rate.

In addition to these times, we consider the corresponding quantities ν_i : for $i = 1, 2, 3, 4, \nu_i$ is the index of the interval (S_{i-1}, S_i) in which the event corresponding to H_i happens; remember that S_i is the time when the *i*th server is created, so that ν_i is the number of servers in the system at time H_i . Note that ν_1 corresponds to Definition 1 (with the same provision between ν_1 and ν^R as between H_1 and T^R). In every simulation, for a fixed N, the averages of the quantities ν_i and H_i are calculated for the value $\rho = 2$ over 10^4 iterations of the system which proved to be sufficient in term of numerical stability. The number of peers N ranges up to 5.10^7 .

Validation of Approximation B. Proposition 2.1 and the subsequent discussion suggest that Approximation B holds at least up to time T^R : in view of the above remark on the difference between T^R and H_1 , we justify that Approximation B holds up to time H_1 by comparing $\mathbb{E}(H_1)$ and $\mathbb{E}(\nu_1)$ to $\mathbb{E}(T^R)$ and $\mathbb{E}(\nu^R)$, respectively. Corollary 5.1 suggests that $\mathbb{E}(\nu_1) \approx A_1 N^{1/(\rho+2)}$ for some constant A_1 , and $\mathbb{E}(H_1) \approx$ $\log(N)/(\rho+2)$. Figure 2 shows the graphs of $\log(\mathbb{E}(\nu_1))$ and $\mathbb{E}(H_1)$ against $\log(N)$: the straight lines depicted prove a good agreement with the theory. Moreover, via a fitting procedure, one can compute the slopes of these lines, and the results are summarized in Table 1 in the row labelled "Min" (each row of the table corresponds to a routing policy, the row "Min" corresponding to the policy considered so far when an incoming peer is router to the less loaded server).

Policy	$ u_1 \mid H_1 $	$ u_2 \mid H_2$	$ u_3 \mid H_3$
Min	$0.248 \ / \ 0.256$	$0.376 \ / \ 0.515$	$0.315 \ / \ 0.329$
Random	$0.247 \ / \ 0.257$	$0.371 \ / \ 0.508$	$0.238 \ / \ 0.253$

TABLE 1. Coefficients of growth rates for the three different Heuristics (1), (2) and (3). For instance, $\nu_2 \approx N^{0.38}$ and $H_2 \approx 0.51 \log N$ under the Min policy.



FIGURE 3. Comparison of $\mathbb{E}(H_1)$, $\mathbb{E}(T^D)$ and $\mathbb{E}(H_2)$ against N.

We see that simulations exhibit a slope of 0.248 for ν_1 and of 0.256 for H_1 , whereas the theory predicts 0.25 in both cases (because $\rho = 2$). This justifies the fact that Approximation B holds at least until H_1 , so that T^R predicted by Corollary 5.1 is a good estimate of H_1 .

Accuracy of Bins and Balls Models. We now analyze the analytical results predicted by the random and deterministic bins and balls models. The time H_2 , defined as the last time when there is no empty server in the system, plays a special role: it appears in Figure 1 that H_2 is closely related to the end of the first regime. Right after H_2 , the number of servers grows sharply, and before H_2 empty servers seem to be negligible. The question is to know which one of $T^R \approx H_1$ and T^D is closer to H_2 .

Not surprisingly, Figure 3 shows that H_1 is significantly different from H_2 . The reason has already been alluded to in Section 2: results obtained for the random model point out a local behavior. The first empty bin in the random model arrives in a region where still many peers arrive in each interval: for instance, although no peer arrives in the first empty interval, it can be shown that a great number of peers arrive in the surrounding intervals. In some sense, many peers should arrive in this interval, but they don't because a server is created very quickly. This is a rare event which does not give insight into the global equilibrium of the system: the order of magnitude $N^{1/(\rho+2)}$ provided by the random bins and balls model is thus misleading.

The deterministic model provides a better answer. In this model, the sizes of bins are not random, and the stochastic fluctuations arising in the random model do not occur. The deterministic model smooths the local behavior that appears in the random model, and the order of magnitude $(N/\log N)^{1/(\rho+1)}$ gives more insight into the global situation of the system. When only a few peers arrive in an interval, it really means that the equilibrium begins to shift. One can check in Figure 3 that the theoretical result T^D predicted in Equation (20) by the deterministic model is closer to H_2 than H_1 .

Although the deterministic model indeed improves the approximation, H_2 still seems much larger that T^D . A first information derived thanks to our bins and balls models is that the first order approximation of the times H_i is logarithmic, whereas the first order approximation for the indexes ν_i is polynomial.

Second, and more interestingly, the deterministic model yields a reasonable estimate of the number of servers at the end of the first regime, roughly equal to ν_2 . Indeed, one can check on Table 1 that simulations give a slope of 0.38 for ν_2 when the deterministic model predicts 0.33. The random model predicts 0.25, so a substantial improvement in accuracy is obtained when using the deterministic model. As will be seen when discussing the two last Heuristics (3) and (4), Approximation B is very likely to hold up to time T^D and not after, which entails the following puzzling observation. Between T^D and T_2 , only few servers are created since $\nu^{D} \approx \nu_{2}$; however to create these servers, a time greater than the time predicted by the branching Assumption B is needed. Indeed, one can see in Table 1 that the coefficient of growth is equal to 0.38 for ν_2 and 0.51 for H_2 , which are significantly different. This means that the branching Assumption B does not hold after time T^D , and in particular that it does not hold until time H_2 . Between T^D and H_2 , it could happen that there is a small fraction of empty servers which has nonetheless a significant impact on the system. A similar phenomenon has been observed in Sanghavi et al. [SHM07].

Two Other Heuristics. We now discuss the two last Heuristics (3) and (4) and conclude that they give the same estimate as T^D , with of course less precise asymptotic.

Recall that H_3 is the first time when a server empties. This time has an appealing motivation: since departures occur uniformly at random from any non-empty server, a departure from a server with only one peer means that such servers represent a significant proportion of servers. Simulations show that ν_3 and H_3 have similar behavior as before (polynomial and logarithmic growths, respectively). Results in Table 1 show that the slope for H_3 is close to the exponent of ν_3 , suggesting that Approximation B holds until H_3 . Moreover, the value 0.31 of the slope is close to the theoretical value 0.33 predicted by the deterministic model, which supports the idea that T^D is as good an estimate as one can do within the scope of Approximation B.

Finally, let us discuss the last and maybe most natural Heuristic (4). Throughout this chapter, we have tried to estimate the time when the equilibrium of the system begins to shift: the time H_4 corresponds to the first time the input rate is smaller than the output rate. More precisely, the input rate i(t) into the system is simply the number of peers that are not active at time t times ρ . Note that the number of asleep peers at time t can be neatly written as

$$\frac{1}{\rho}i(t) = \sum_{i=1}^{N} \mathbb{1}_{\{E_i^{\rho} > t\}}$$

if E_i^{ρ} is the time at which the *i*th peer awakes; the $(E_i^{\rho}, i \geq 1)$ are then i.i.d. with common distribution the exponential distribution with parameter ρ . On the other hand, as long as Approximation B holds, the output rate o(t) is just the number of servers at time *t* (since peers require a service of mean one). Initially, $i(0) = \rho N$ and o(0) = 1, and eventually $i(\infty) = 0$ and $o(\infty) = N$. To study the time at which the equilibrium of the system begins to shift, it is natural to consider the first time H_4 at which i(t) < o(t). As shown in the following, this leads to the order of magnitude T^D given by the deterministic model.

Assuming that Approximation B holds for times $t < H_4$ — note that in contrast with our heuristic, this assumption is not easy to justify — the problem can be cast in terms of the random bins and balls problem. Let Z_N^x be the number of balls that fall in the x first intervals:

$$Z_N^x = \sum_{i=1}^x \eta_i(N) = \sum_{i=1}^N \mathbb{1}_{\{E_i^{\rho} \le T_x\}},$$

where T_x is given by Equation (1) for any x > 0. The index ν_4 then corresponds to

$$\nu_4 = \inf \left\{ x : N - Z_N^x \stackrel{\text{def.}}{=} \widetilde{Z}_N^x < \frac{x}{\rho} \right\}.$$

The asymptotic behavior of $\mathbb{E}(\widetilde{Z}_N^x)$ when x goes to infinity with N is easy to derive:

$$\mathbb{E}(\widetilde{Z}_N^x) = N \sum_{i>x+1} \mathbb{E}P_i \sim \alpha N \sum_{i>x+1} i^{-\rho-1} \sim \frac{\alpha}{\rho} N x^{-\rho}.$$

Therefore $\mathbb{E}(\widetilde{Z}_N^x) \approx x$ for $x \approx N^{1/(\rho+1)}$, i.e., ν_4 is of order $N^{1/(\rho+1)}$, which is essentially the same order of magnitude as in the deterministic model. Although ν_4 has been cast in terms of a random bins and balls problem, there is, in contrast with ν^R and ν^D , no discrepancy to be expected between the random and the deterministic models. Indeed, \widetilde{Z}_N^x depends on the cumulative number of balls in the last intervals, and not on the number of balls in some interval. Thus a law of large numbers effect will prevail and \widetilde{Z}_N^x will indeed behave as its first moment, i.e., results of the random and deterministic models will be the same. Rigorous mathematical analysis could be done to prove this result, but in our view, considering H_1 has one main advantage: Proposition 2.1 is almost a rigorous justification of Approximation B. When considering H_4 and more generally any other time, we were not able to provide such a strong justification. And as we have seen in the case of H_2 , Approximation B does not hold for the whole first regime, and a strong justification as Proposition 2.1 is therefore valuable.

In conclusion, we have provided numerous arguments showing that Approximation B holds until times of order of $N^{1/(\rho+1)}$, which corresponds to T^D , H_4 and at least one particular set of simulations — H_3 ; the deterministic model moreover gives a reasonable estimate on the *number* of servers at the end of the first regime. Since finally Proposition 2.1 gives a strong justification of the heuristic underlying the deterministic model, this technical approach is very valuable.



FIGURE 4. Comparison of Min (dotted) and Random (solid) for the Heuristic (1) corresponding to H_1 and ν_1 .



FIGURE 5. Comparison of Min (dotted) and Random (solid) for the Heuristic (2) corresponding to H_2 and ν_2 .

The branching Approximation B however does not hold until H_2 , which seems to best approximate the end of the first regime: between T^D and H_2 , only few servers are created, which nonetheless amount to a time much larger than the time predicted by (1). Since Approximation B is the cornerstone of our analysis, different tools should be used to approximate H_2 .

To conclude this chapter, we discuss a possible modification of the system as well as some natural extensions.

Comparison of Routing Policies. In the system considered so far, an incoming peer is queued at the server with the least number of queued peers. From a practical viewpoint it requires the knowledge of the state of every server, which is impractical. A simpler solution, discussed now, consists in routing a peer to a server chosen uniformly at random. We discuss through simulations the effect of this new policy on the different Heuristics (1), (2) and (3); numerical results are summarized in Table 1 in the row labelled "Random" and in Figures 4, 5 and 6.



FIGURE 6. Comparison of Min (dotted) and Random (solid) for the Heuristic (3) corresponding to H_3 and ν_3 .

First of all we observe that Heuristic (1) is insensitive to the policy: Table 1 shows that the growth rates are not affected by the policy, and the two curves corresponding to the Min or Random policy are almost undistinguishable on Figure 4.

A similar observation goes with Heuristic (2): the time H_2 is essentially unaffected, and the number of servers ν_2 is only affected through a change in constant, i.e., $\nu_2 \approx A_m N^{0.37}$ under the Min policy and $\nu_2 \approx A_r N^{0.37}$ under the random one, with two different constants $A_m \neq A_r$. Thus for our purposes, the routing policy seems to have no effect on the length of the first regime.

An interesting exception concerns Heuristic (3): Table 1 shows that for the first time H_3 when a server becomes empty, the policy has a great influence. This is easily understandable: under the Min policy, it is much harder for a server to become empty, because least loaded servers are selected by incoming peers.

Extensions. We conclude this chapter with some possible extensions. The first one consists in allowing peers to leave the system instead of becoming a server, once downloading the file. This effect is called "free riding", it has a significant impact on real peer-to-peer systems. Our results extend to cover this case: the output rate μ of a server corresponds to the rate at which a peer finishes its download. Hence $p\mu$ corresponds to the rate at which a server creates another new server, if $0 \le p \le 1$ is the probability for a peer to stay in the system once downloading the file. Thus our results naturally extend by just replacing ρ by ρ/p . On the other hand the problem of peers who leave while being a server is not that easy to handle, since one would have to deal with the peers it was serving at the time it disconnects. This is moreover a question which is more relevant in a stationary context, with a continuous steady flow of incoming peers, and where one is interested in stability properties such as in the previous chapter.

Another extension concerns the download time, which was supposed to be exponentially distributed. If the service discipline is FIFO, then these results could probably be extended to an arbitrary service distribution. Indeed, there would still be a period of time during which the population of servers resembles a branching process, which is no longer a Yule process. It then becomes a general binary Bellman-Harris process, and the sequence of split times has a structure which closely resembles to (1) asymptotically, see for instance Athreya [AK77, Ath69] and Härnqvist [H**B**1]. Nonetheless there is an additional technical difficulty, which we were not able to overcome, see the discussion in Section IV.6. Note that if one wants to study arbitrary service requirements under the Processor-Sharing discipline, then results on branching processes become useless, since the departure process from an overloaded Processor-Sharing queue is in general not easy to describe.

References of Chapter III

- [AK77] K. B. Athreya and Niels Keiding. Estimation theory for continuoustime branching processes. Sankhya: The Indian Journal of Statistics, 89(A):101–123, 1977.
- [AN72] K. B. Athreya and P. E. Ney. *Branching processes*. Springer, 1972.
- [Asm87] Søren Asmussen. Applied probability and queues. John Wiley & Sons, Chichester, 1987.
- [Ath69] Krishna B. Athreya. On the supercritical one dimensional age dependent branching processes. Annals of Mathematical Statistics, 40(3):743–763, June 1969.
- [BHJ92] A. D. Barbour, Lars Holst, and Svante Janson. Poisson approximation. The Clarendon Press Oxford University Press, New York, 1992.
- [CF76] Endre Csáki and Antónia Földes. On the first empty cell. *Studia Scientarum Mathematicarum Hungarica*, 11:373–382, 1976.
- [CN04] F. Clévenot and P. Nain. A simple fluid model for the analysis of the squirrel peer-to-peer caching system. In *Proceedings of IEEE INFOCOM* 2004, volume 1, pages 86–95, Hong Kong, March 2004.
- [GFS⁺03] Z. Ge, D. R. Figueiredo, Jaiswal Sharad, J. Kurose, and D. Towsley. Modeling peer-peer file sharing systems. In *Proceedings of IEEE INFO-COM 2003*, volume 3, pages 2188–2198, San Francisco, CA, USA, March 2003.
- [GHP07] Alexander V. Gnedin, Ben Hansen, and Jim Pitman. Notes on the occupancy problem with infinitely many boxes: General asymptotics and power laws. *Probability Surveys*, 4:146–171, 2007.
- [H81] Martin Härnqvist. Limit theorems for point processes generated in a general branching process. Advances in Applied Probability, 13:650–668, 1981.
- [HJ08] Hsien-Kuei Hwang and Svante Janson. Local limit theorems for finite and infinite urn models. *Annals of Probability*, 36(3):992–1022, 2008.
- [Kin78] J. F. C. Kingman. Random partitions in population genetics. Proceedings of the Royal Society. London. Series A. Mathematical, Physical and Engineering Sciences, 361(1704):1–20, 1978.
- [MV05] Laurent Massoulié and Milan Vojnović. Coupon replication systems. In Proceedings of SIGMETRICS'05, volume 33, pages 2–13, Banff, Alberta, Canada, June 2005.
- [QS04] Dongyu Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In *Proceedings of SIGCOMM'04*, pages 367–378, New York, NY, USA, August 2004. ACM.
- [RS09] Philippe Robert and Florian Simatos. Occupancy schemes associated to Yule processes. Advances in Applied Probability, 41(2):600–622, June 2009.
- [SHM07] S. Sanghavi, B. Hajek, and L. Massoulié. Gossiping with Multiple Messages. In *Proceedings of IEEE INFOCOM 2007*, pages 2135–2143, May 2007.
- [Wil91] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.
[YdV06] Xiangying Yang and Gustavo de Veciana. Performance of peer-to-peer networks: service capacity and role of resource sharing policies. *Perfor*mance Evaluations, 63(3):175–194, 2006.

CHAPTER IV

Occupancy Schemes Associated to Yule Processes

Contents

1. Introduction	135
2. A Bins and Balls Problem in Random Environment	137
3. Convergence of Point Processes	138
4. Asymptotic Behavior of the Indices of the First Empty	
Bins	145
5. Rare Events	148
6. Generalizations	152
7. Comparison with the Renewal Case	154
Appendix A. Proof of Proposition 5.1	158
References of Chapter IV	

1. Introduction

Occupancy schemes in terms of bins and balls offer a very flexible and elegant way to formulate various problems in computer science, biology and applied mathematics for example. One of the earliest models investigated in the literature consists in throwing m balls at random into n identical bins. Asymptotic behavior of occupancy variables have been analyzed when n grows to infinity, with different scalings in n for the variable m. The books by Johnson and Kotz [JK77] and Kolchin *et al.* [KS78] are classical references on this topic. See also Chapter 6 of Barbour *et al.* [BHJ92] for a recent presentation of these problems.

An extension of these models is when there is an infinite number of bins and a probability vector (p_n) on \mathbb{N} describing the way balls are sent: for $n \ge 0$, p_n is the probability that a ball is sent into the *n*th bin. In one of the first studies in this setting, Karlin [**Kar67**] analyzed the asymptotic behavior of the number of occupied bins. More recently Hwang and Janson [**HJ08**] proves in a quite general framework central limit results for these quantities. In this setting, some additional variables are also of interest like the sets of indices of occupied or empty bins, adding a geometric component to these problems. For specific probability vectors (p_n) Csáki and Földes [**CF76**] and Flajolet and Martin [**FM85**] investigated the index of the first empty bin. See the recent survey Gnedin *et al.* [**GHP07**] for more references on the occupancy problem with infinitely many bins.

A further extension of these stochastic models consists in considering random probability vectors. Gnedin [Gne04] (and subsequent papers) analyzed the case where (p_n) decays geometrically fast according to some random variables, i.e., for $n \geq 1$, $p_n = \prod_{i=1}^{n-1} Y_i(1-Y_n)$ where (Y_i) are i.i.d. random variables on (0,1).

Various asymptotic results on the number of occupied bins in this case have been obtained. The random vector can be seen as a "random environment" for the bins and balls problem, it complicates significantly the asymptotic results in some cases. In particular, the indices of the bins in which the balls fall are no longer independent random variables as in the deterministic case.

The general goal of this chapter is to investigate in detail the impact of this randomness for a bins and balls problem associated to a Yule process, see Athreya and Ney [AN72] for the definition of a Yule process. This (quite natural) stochastic model has its origin in network modeling, see Simatos *et al.* [SRG08] for a detailed presentation. It can be described as follows: the non-decreasing sequence (t_n) of split times of the Yule process defines the bins, the *n*th bin, $n \ge 1$, being the interval $(t_{n-1}, t_n]$. The locations of balls are represented by independent exponential random variables with parameter ρ . The main problem investigated here concerns the asymptotic description of the set of indices of first empty bins when the number of balls goes to infinity. Mathematically, it is formulated as a convergence in distribution of rescaled point processes having Dirac masses at the indices of empty bins.

For $n \ge 1$, if P_n is the probability that a ball falls into the *n*th bin, it is easily seen that, for a large n, P_n has a power law decay, it can be expressed as $VE_n/n^{\rho+1}$ where (E_n) are i.i.d. exponential random variables with parameter 1 and V some independent random variable related to the limit of a martingale. The randomness of the probability vector (P_n) has two components: one which is a part of an i.i.d. sequence, changing from one bin to another, and the other being "fixed once for all" inducing a dependency structure. As it will be seen, the two components have separately a significant impact on the qualitative behavior of this model.

Convergence in Distribution and Rare Events. Because the variables (E_n) can be arbitrarily small with positive probability, empty bins are likely to be created earlier (i.e., with smaller indices) than for a deterministic probability vector with the same power law decay. It is shown in fact that, for the convergence in distribution, the first empty bins occur around indices of the order of $n^{1/(\rho+2)}$ instead of $(n/\log n)^{1/(1+\rho)}$ in the deterministic case.

The variable V has a more subtle impact, when $\rho > 1$ it is shown that, due to some heavy tail property of V^{-1} , rare events affect the asymptotic behavior of *averages* of some of the characteristics. For $\alpha \in [1/(2\rho + 1), 1/(\rho + 2))$, despite that the number of empty bins with indices of order n^{α} converges in distribution to 0, the corresponding average converges to $+\infty$. When $\rho < 1$, the average is converging to 0 for any $\alpha < 1/(\rho+2)$. A phase transition phenomenon at $\rho = 1$ has been identified through simulations in a related context, communication networks, in Saddi and Guillemin [SG07]. It is not apparent as long as convergence in distribution is concerned but it shows up when average quantities are considered. This phenomenon is due to rare events related to the total size of the $\lfloor \rho \rfloor$ first bins: On these events, the indices of the first empty bins are of the order $n^{1/(2\rho+1)} \ll n^{1/(\rho+2)}$ and a lot them are created at this occasion. See Proposition 5.3 and Corollary 5.1 for a precise statement of this result. Concerning the generality of the results obtained, it is believed that some of them hold in a more general setting, for the underlying branching process for example, see Section 6. **Point Processes.** Technically, one mainly uses point processes on \mathbb{R}_+ to describe the asymptotic behavior of the indices of the first empty bins and not only the index of the first one (or the subsequent ones) as it is usually the case in the literature. It turns out that it is quite appropriate in our setting to get a full description of the set of the first empty bins and, moreover, it reduces the technicalities of some of the proofs. One of the arguments for the proofs of the convergence results is a simple convergence result of two-dimensional point processes to a Poisson point process with some intensity measure. A one-dimensional equivalent of this point of view is implicit in most of the papers of the literature, see Hwang and Janson [HJ08] and Gnedin *et al.* [GIR08] in particular.

The chapter is organized as follows. Section 2 introduces the stochastic model investigated. The main results concerning convergence of related point processes in \mathbb{R}^2_+ are presented in Section 3. Convergence results for the indices of empty bins are proved in Section 4. Section 5 investigates in detail the case $\rho \geq 1$. Section 6 presents some possible extensions and Section 7 compares our results with a related model investigated by Gnedin *et al.*. The Appendix A is devoted to the proof of some technical estimates.

2. A Bins and Balls Problem in Random Environment

The stochastic model is described in detail and some notations are introduced. The following problem was already introduced in Chapter III with slightly different notations. In order for this chapter to be self-contained, some simple results already proved in Chapter III are proved again here.

The Bins. Let (E_i) be a sequence of i.i.d. exponential random variables with parameter 1. Define the non-decreasing sequence (t_n) by, $t_0 = 0$ and, for $n \ge 1$,

$$t_n = \sum_{i=1}^n \frac{1}{i} E_i.$$

It is easy to check that for $x \ge 0$,

(1)
$$\mathbb{P}(t_n \le x) = \mathbb{P}(\max(E_1, E_2, \dots, E_n) \le x) = (1 - e^{-x})^n.$$

The *n*th bin will be identified by the interval $(t_{n-1}, t_n]$.

If $H_n = 1 + 1/2 + \cdots + 1/n$ is the *n*th harmonic number, since $(t_n - H_n)$ is a square integrable martingale whose increasing process is given by

$$\mathbb{E}\left((t_n - H_n)^2\right) = \sum_{i=1}^n \frac{1}{i^2},$$

then $(M_n) \stackrel{\text{def.}}{=} (t_n - \log(n+1))$ is almost surely converging to some finite random variable M_{∞} . See Neveu [Nev72] or Williams [Wil91]. By using Equation (1), it is not difficult to get that the distribution of M_{∞} is given by

(2)
$$\mathbb{P}(M_{\infty} \le x) = \exp\left(-e^{-x}\right), \quad x \in \mathbb{R}.$$

An alternative description of the sequence (t_n) is provided by the split times of a Yule (branching) process starting with one individual. See Athreya and Ney [**AN72**]. **The Balls.** The locations of the balls are given by an independent sequence (B_j) of i.i.d. exponential random variables with parameter ρ for some $\rho > 0$.

Conditionally on the point process (t_n) associated with the location of bins, the probability that a given ball falls into the *n*th bin $(t_{n-1}, t_n]$ is given by

$$P_n = \mathbb{P}\left[B_1 \in (t_{n-1}, t_n] \, \Big| (t_n)\right] = e^{-\rho t_{n-1}} - e^{-\rho t_n} = e^{-\rho t_{n-1}} \left(1 - e^{-\rho E_n/n}\right).$$

This quantity can be rewritten as

(3)
$$P_n = \frac{1}{n^{\rho+1}} W_n^{\rho} Z_n$$
, with $Z_n = n \left(1 - e^{-\rho E_n/n} \right)$ and $W_n = e^{-M_{n-1}}$.

The variables W_n and Z_n are independent random variables with different behavior.

- (1) The variables (Z_n) are independent and converge in distribution to an exponentially distributed random variable with parameter $1/\rho$.
- (2) The random variables (W_n) converge almost surely to the finite random variable $W_{\infty} = \exp(-M_{\infty})$ which is exponentially distributed with parameter 1.

This suggests an asymptotic representation of the sequence (P_n) as

(4)
$$\left(\frac{1}{n^{\rho+1}}W^{\rho}_{\infty}F_n\right),$$

where (F_n) is an i.i.d. sequence of exponential random variables with mean ρ independent of W_{∞} . The sequence (P_n) has a power law decay with a random coefficient consisting of the product of two terms: a fixed random variable W_{∞}^{ρ} and the other being an element of an i.i.d. sequence. As it will be seen, these two terms have a significant impact on the bins and balls problem studied in this paper.

3. Convergence of Point Processes

One of the main result, Theorem 4.1 in the next section, which establishes convergence results for the indices of the first empty bins is closely related to the asymptotic behavior of the point process $\{(i/n^{1/(2+\rho)}, nP_i), i \ge 1\}$ on \mathbb{R}^2_+ . For this reason, some results on convergence of point processes in \mathbb{R}^2_+ are first proved. The point process associated to the sequence (nP_i) appears quite naturally, especially in view of the Poisson transform used in the proof of Theorem 4.1. This is also a central variable in Hwang and Janson [**HJ08**] in some cases.

An important tool to study point processes in \mathbb{R}^d_+ for some $d \ge 1$ is the Laplace transform: If $\mathcal{N} = \{t_n, n \ge 1\}$ is a point process and f a function in $C_c^+(\mathbb{R}^d_+)$, the set of non-negative continuous functions with a compact support, it is defined as $\mathbb{E}(\exp(-\mathcal{N}(f)))$, where

$$\mathcal{N}(f) \stackrel{\text{def.}}{=} -\sum_{n\geq 1} f(t_n).$$

This functional uniquely determines the distribution of \mathcal{N} and it is a key tool to establish convergence results. See Neveu [Nev77] and Dawson [Daw93] for a comprehensive presentation of these questions. In the following, the quantity $\mathcal{N}(A)$ denotes the number of t_n 's in the subset A of \mathbb{R}^d_+ .

The main results of this section establish convergence in distribution to mixed Poisson point processes, i.e., distributed as a Poisson point process with a parameter which is a random variable. A natural tool in this domain is the Chen-Stein approach which gives the convergence in distribution and, generally, quite good bounds on the convergence rate. See Chapter 10 of Barbour *et al.* [**BHJ92**] for example. This has been used in the previous chapter when the probability vector is deterministic. For some of the results of this section, this approach can probably also be used. Unfortunately, due to the almost surely converging sequence (W_n) creating a dependency structure, it does not seem that the main convergence result, Theorem 3.1, can be proved in a simple way by using Chen-Stein's method. The main problem being of conditioning on the variable W_{∞} and keeping at the same time upper bounds on the total variation distance converging to 0.

Condition C. A sequence of independent random variables (X_i) satisfies Condition C if there exist some $\alpha > 0$ and $\eta > 0$ such that, for all $i \ge 1$,

(5)
$$|\mathbb{P}(X_i \le x) - \alpha x| \le Cx^2, \text{ when } 0 \le x \le \eta.$$

The following proposition is a preliminary result that will be used to prove the main convergence results for the indices of the first empty bins.

PROPOSITION 3.1 (Convergence to a Poisson process). For $\delta > 0$ and $n \ge 1$, let \mathcal{P}_n be the point process on \mathbb{R}^2_+ defined by

$$\mathcal{P}_n \stackrel{\text{def.}}{=} \left\{ \left(\frac{i}{n^{1/(\delta+1)}}, \frac{n}{i^{\delta}} X_i \right), i \ge 1 \right\},\,$$

where (X_i) is a sequence of non-negative independent random variables satisfying Condition C. Then the sequence of point processes (\mathcal{P}_n) converges in distribution to a Poisson point process \mathcal{P} in \mathbb{R}^2_+ with intensity measure $x^{\delta} dx dy$ on \mathbb{R}^2_+ . In particular, its Laplace transform is given by

(6)
$$\mathbb{E}(\exp[-\mathcal{P}(f)]) = \exp\left(-\alpha \int_{\mathbb{R}^2_+} \left(1 - e^{-f(x,y)}\right) x^{\delta} \, dx \, dy\right), \quad f \in C^+_c(\mathbb{R}^2_+).$$

See Robert [**Rob03**] for the definition and the main properties of Poisson processes in general state spaces.

PROOF. There exists some $\eta_0 > 0$ such that $\mathbb{P}(X_i \leq x) \leq 2\alpha x$ for $0 \leq x \leq \eta_0$ and all $i \geq 1$. Let $f \in C_c^+(\mathbb{R}^2_+)$ be such that f is differentiable with respect to the second variable. There is some K > 0 so that the support of f is included in $[0, K] \times [0, K]$, define $g(x, y) = 1 - \exp(-f(x, y))$, then by independence of the variables $X_i, i \geq 1$,

$$\log \mathbb{E}\left(e^{-\mathcal{P}_n(f)}\right) = \sum_{i=1}^{+\infty} \log\left(1 - \mathbb{E}\left[g\left(\frac{i}{n^{1/(\delta+1)}}, \frac{n}{i^{\delta}}X_i\right)\right]\right).$$

Since

$$\mathbb{E}\left[g\left(\frac{i}{n^{1/(\delta+1)}},\frac{n}{i^{\delta}}X_i\right)\right] \leq \mathbb{P}\left(X_i \leq K\frac{i^{\delta}}{n}\right) \mathbb{1}_{\{i \leq Kn^{1/(\delta+1)}\}},$$

the elementary inequality $|\log(1-y) + y| \le 3y^2/2$ valid for $0 \le y \le 1/2$ shows that there exists some $n_0 \ge 1$ such that

$$\left| \log \mathbb{E} \left(e^{-\mathcal{P}_n(f)} \right) + \sum_{i=1}^{+\infty} \mathbb{E} \left[g \left(\frac{i}{n^{1/(\delta+1)}}, \frac{n}{i^{\delta}} X_i \right) \right] \right|$$

$$\leq \frac{6(\alpha K)^2}{n^2} \sum_{i=1}^{\lfloor K n^{1/(\delta+1)} \rfloor} i^{2\delta} \leq 6\alpha^2 K^{2\delta+3} \frac{1}{n^{1/(\delta+1)}}$$

holds for $n \ge n_0$. It is therefore enough to study the asymptotics of the series of the left hand side of the above inequality. For $x \ge 0$, by using Fubini's Theorem, one gets

$$\mathbb{E}\left(g\left(x,\frac{n}{i^{\delta}}X_{i}\right)\right) = -\int_{0}^{+\infty}\frac{\partial g}{\partial y}(x,y)\mathbb{P}\left(X_{i} \leq yi^{\delta}/n\right)\,dy$$

By using again Condition C, one obtains that the log of the Laplace transform of \mathcal{P}_n has the same asymptotic behavior as

$$-\alpha \frac{1}{n^{1/(\delta+1)}} \sum_{i=1}^{+\infty} \int_0^{+\infty} \frac{\partial g}{\partial y} \left(\frac{i}{n^{1/(\delta+1)}}, y\right) y\left(\frac{i}{n^{1/(\delta+1)}}\right)^{\delta} dy$$

which is a Riemann sum converging to

$$-\alpha \int_{\mathbb{R}^2_+} \frac{\partial g}{\partial y}(x, y) y x^{\delta} \, dx dy = \alpha \int_{\mathbb{R}^2_+} \left(1 - e^{-f(x, y)} \right) x^{\delta} \, dx dy.$$

This shows in particular that for any compact set H of \mathbb{R}^2_+ , then

$$\sup_{n\geq 1} \mathbb{E}(\mathcal{P}_n(H)) < +\infty$$

the sequence (\mathcal{P}_n) is therefore tight for the weak topology, see Dawson [Daw93].

By the convergence result, if \mathcal{P} is any limiting point of the sequence (\mathcal{P}_n) , for any function $f \in C_c^+(\mathbb{R}^2_+)$ such that $y \mapsto f(x, y)$ is differentiable, then the Laplace transform of \mathcal{P} at f is given by the right hand side of Equation (6). By density of these functions f in $C_c^+(\mathbb{R}^2_+)$ for the uniform topology, this implies that \mathcal{P} is indeed a Poisson point process with intensity measure $x^{\delta} dx dy$ on \mathbb{R}^2_+ . The proposition is proved.

The above result can be (roughly) restated as follows: for the indices of the order of $n^{1/(\delta+1)}$, the points nX_i/i^{δ} lying in some finite fixed interval converge to an homogeneous Poisson point process. The following proposition gives an asymptotic description of the indices of the points nX_i/i^{δ} but for indices of the order of n^{κ} with $1/(\delta+1) < \kappa < 1/\delta$. It shows that, on finite intervals, these points accumulate at rate $n^{(1+\delta)\kappa-1}$ according to the Lebesgue measure with some density.

PROPOSITION 3.2 (Law of Large Numbers). If, for $1/(1 + \delta) < \kappa < 1/\delta$ and for $n \ge 1$, \mathcal{P}_n^{κ} is the point process on \mathbb{R}^2_+ defined by

$$\mathcal{P}_n^{\kappa}(f) = \frac{1}{n^{(1+\delta)\kappa-1}} \sum_{i=1}^{+\infty} f\left(\frac{i}{n^{\kappa}}, \frac{n}{i^{\delta}} X_i\right), \quad f \in C_c^+(\mathbb{R}^2_+),$$

where (X_i) is a sequence of non-negative independent random variables satisfying Condition C, then the sequence (\mathcal{P}_n^{κ}) converges in distribution to the deterministic

.

measure $\mathcal{P}^{\kappa}_{\infty}$ defined by

$$\mathcal{P}^{\kappa}_{\infty}(f) = \alpha \int_{\mathbb{R}^2_+} f(x, y) x^{\delta} \, dx dy, \quad f \in C^+_c(\mathbb{R}^2_+).$$

PROOF. Let $f \in C_c^+(\mathbb{R}^2_+)$ be such that f is differentiable with respect to the second variable. As before, the convergence result is proved for such a function f, the generalization to an arbitrary function $f \in C_c^+(\mathbb{R}^2_+)$ follows the same lines as the previous proof (relative compactness argument and identification of the limit). Let K > 0 such that the support of f is included in $[0, K] \times [0, K]$. One has

$$\mathbb{E}\left(\mathcal{P}_{n}^{\kappa}(f)\right) = -\frac{1}{n^{(1+\delta)\kappa-1}} \sum_{i=1}^{+\infty} \int_{0}^{+\infty} \frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right) \mathbb{P}(X_{i} \leq yi^{\delta}/n) \, dy,$$

as in the previous proof, by using Condition (5) and the fact that $\kappa < 1/\delta$, one gets the equivalence

$$\mathbb{E}\left(\mathcal{P}_{n}^{\kappa}(f)\right) \sim -\alpha \frac{1}{n^{\kappa}} \sum_{i=1}^{+\infty} \int_{0}^{+\infty} \frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right) y\left(\frac{i}{n^{\kappa}}\right)^{\delta} \, dy,$$

therefore,

$$\lim_{n \to +\infty} \mathbb{E}\left(\mathcal{P}_n^{\kappa}(f)\right) = -\alpha \int_0^{+\infty} \int_0^{+\infty} \frac{\partial f}{\partial y}\left(x, y\right) y x^{\delta} \, dx dy = \alpha \int_{\mathbb{R}^2_+} f(x, y) x^{\delta} \, dx dy.$$

By independence of the X_i 's the second moment of the difference

$$\mathcal{P}_{n}^{\kappa}(f) - \mathbb{E}\left(\mathcal{P}_{n}^{\kappa}(f)\right) \\ = -\frac{1}{n^{(1+\delta)\kappa-1}} \sum_{i=1}^{+\infty} \int_{0}^{+\infty} \frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right) \left[\mathbbm{1}_{\{X_{i} \leq yi^{\delta}/n\}} - \mathbb{P}(X_{i} \leq yi^{\delta}/n)\right] \, dy,$$

can be expressed as

$$n^{2((1+\delta)\kappa-1)} \times \mathbb{E}\left(\left[\mathcal{P}_{n}^{\kappa}(f) - \mathbb{E}\left(\mathcal{P}_{n}^{\kappa}(f)\right)\right]^{2}\right)$$

$$= \sum_{i=1}^{+\infty} \mathbb{E}\left(\left[\int_{0}^{+\infty} \frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right) \left[\mathbbm{1}_{\{X_{i} \leq yi^{\delta}/n\}} - \mathbb{P}(X_{i} \leq yi^{\delta}/n)\right] dy\right]^{2}\right)$$

$$\leq K \sum_{i=1}^{+\infty} \int_{0}^{+\infty} \left[\frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right)\right]^{2} \mathbb{E}\left(\left[\mathbbm{1}_{\{X_{i} \leq yi^{\delta}/n\}} - \mathbb{P}(X_{i} \leq yi^{\delta}/n)\right]^{2}\right) dy$$

$$\leq K \sum_{i=1}^{+\infty} \int_{0}^{+\infty} \left[\frac{\partial f}{\partial y}\left(\frac{i}{n^{\kappa}}, y\right)\right]^{2} \mathbb{P}(X_{i} \leq yi^{\delta}/n) dy,$$

by Cauchy-Shwartz's Inequality. The last term is, with the same arguments as for the asymptotics of $\mathbb{E}(\mathcal{P}_n^{\kappa}(f))$, equivalent to

$$Kn^{(1+\delta)\kappa-1} \times \int_{\mathbb{R}^2_+} \left[\frac{\partial f}{\partial y}(x,y)\right]^2 yx^{\delta} dxdy.$$

In particular, the sequence $(\mathcal{P}_n^{\kappa}(f))$ converges in L_2 (and therefore in distribution) to $\mathcal{P}_{\infty}^{\kappa}(f)$. The proposition is proved.

The main convergence result can now be established.

THEOREM 3.1. If, for $n \ge 1$, \mathcal{P}_n is the point process on \mathbb{R}^2_+ defined by

$$\mathcal{P}_n = \left\{ \left(\frac{i}{n^{1/(\rho+2)}}, nP_i\right), i \ge 1 \right\},\$$

where the random variables P_i , $i \ge 1$, are defined by Equation (3), then the sequence (\mathcal{P}_n) converges in distribution and the relation

(7)
$$\lim_{n \to +\infty} \mathbb{E}\left(e^{-\mathcal{P}_n(f)}\right) = \mathbb{E}\left[\exp\left(-\frac{W_{\infty}^{-\rho}}{\rho}\int_{\mathbb{R}^2_+} \left(1 - e^{-f(x,y)}\right)x^{\rho+1}\,dx\,dy\right)\right]$$

holds for any $f \in C_c^+(\mathbb{R}^2_+)$.

In other words the point process \mathcal{P}_n converges in distribution to a mixed Poisson point process: conditionally on W_{∞} , it is a Poisson process with intensity measure $W_{\infty}^{-\rho} x^{\rho+1} dx dy/\rho$.

PROOF. The proof proceeds in several steps. The main objective of these steps is to decouple the sequences (W_i) and (Z_i) defining the (P_i) and then to apply Proposition 3.1.

Step 1. One defines the sequences

$$P_i^1 = \frac{1}{i^{\rho+1}} \widetilde{W}_{\infty}^{\rho} \widetilde{Z}_i, \quad i \ge 1, \quad P_i^2 = \frac{1}{i^{\rho+1}} \widetilde{W}_{\beta_n}^{\rho} \widetilde{Z}_i, \quad i \ge 1,$$

where (β_n) is some sequence of integers converging to $+\infty$. Note that (P_i^2) depends on n. The sequences of random variables $(\widetilde{W}_i, 1 \leq i \leq +\infty)$ and (\widetilde{Z}_i) are assumed to be independent and to have, respectively, the same distribution as $(W_i, 1 \leq i \leq +\infty)$ and (Z_i) defined by Equation (3). Recall that the sequence (\widetilde{W}_i) converges almost surely to \widetilde{W}_{∞} . These sequences define point processes in the following way, for j = 1 and 2,

$$\mathcal{P}_n^j = \left\{ \left(\frac{i}{n^{1/(\rho+2)}}, nP_i^j\right), i \ge 1 \right\}.$$

If f is a non-negative continuous function with compact support on \mathbb{R}^2_+ , because, conditionally on \widetilde{W}_{∞} , the variables $(\widetilde{W}_{\infty}Z_i)$ satisfy Condition C with $\alpha = \widetilde{W}_{\infty}^{-\rho}/\rho$, Proposition 3.1, with $\delta = \rho + 1$, shows that

$$\lim_{n \to +\infty} \mathbb{E}\left(e^{-\mathcal{P}_n^1(f)} \middle| \widetilde{W}_{\infty}\right) = \exp\left(-\frac{\widetilde{W}_{\infty}^{-\rho}}{\rho} \int_{\mathbb{R}^2_+} \left(1 - e^{-f(x,y)}\right) x^{\rho+1} \, dx \, dy\right).$$

Because of the boundedness of these quantities, by Lebesgue's Theorem, the same result holds for the expected values. Therefore, the sequence (\mathcal{P}_n^1) converges in distribution to the point process \mathcal{P} on \mathbb{R}^2_+ whose Laplace transform is given by Equation (7).

Let $K \ge 2$ be such that the support of f is a subset of $[0, K]^2$ and $\varepsilon > 0$. Since the limiting point process \mathcal{P} is almost surely a Radon measure, there exists some $m \in \mathbb{N}$ such that $\mathbb{P}(\mathcal{P}_n^1([0, 2K]^2) \ge m) \le \varepsilon$ for all $n \ge 1$. By uniform continuity, there exists $0 < \eta < 1/2$ such that $|f(u) - f(v)| \le \varepsilon/m$ for $u, v \in \mathbb{R}^2_+$ such that $||u - v|| \le \eta$. For $n \ge 1$, if

$$\mathcal{A} \stackrel{\text{def.}}{=} \{ |\widetilde{W}_{\beta_n}^{\rho} / \widetilde{W}_{\infty}^{\rho} - 1| \ge \eta/2K \} \cup \{ \mathcal{P}_n^1([0, 2K]^2) \ge m \}$$

then

$$\begin{split} \left| \mathbb{E} \left(\exp \left[-\mathcal{P}_n^2(f) \right] \right) - \mathbb{E} \left(\exp \left[-\mathcal{P}_n^1(f) \right] \right) \right| &\leq \mathbb{P}(\mathcal{A}) \\ + \mathbb{E} \left(\left(\exp \left[\sum_{i \geq 1} \left| f \left(\frac{i}{n^{1/(\rho+2)}}, \frac{\widetilde{W}_{\beta_n}^{\rho}}{\widetilde{W}_{\infty}^{\rho}} n P_i^1 \right) - f \left(\frac{i}{n^{1/(\rho+2)}}, n P_i^1 \right) \right| \right] - 1 \right) \mathbb{1}_{\mathcal{A}^c} \right) \\ &\leq \mathbb{P}(|W_{\beta_n}^{\rho}/W_{\infty}^{\rho} - 1| \geq \eta/2K) + (1 + e^{\varepsilon})\varepsilon, \end{split}$$

hence, by the almost sure convergence of (W_n) to W_{∞} , the right hand side of the last relation can be arbitrarily small as n goes to infinity. One concludes that the sequence (\mathcal{P}_n^2) also converges in distribution to the point process \mathcal{P} .

Step 2. For $n \ge 1$, define

$$\beta_n = \left\lfloor n^{1/(\rho+2)} / \log n \right\rfloor$$

then it will be shown that the point processes

$$\mathcal{Q}_n = \left\{ \left(\frac{i}{n^{1/(\rho+2)}}, nP_i \right), 1 \le i \le \beta_n \right\}$$

converge to the measure identically null. It is sufficient to prove that for any $f \in C_c^+(\mathbb{R}_+)$, the sequence $(\mathcal{Q}_n(f))$ converges in distribution to 0. For a fixed *i*, the sequence (nP_i) converges in distribution to infinity, since *f* is continuous with compact support and therefore bounded, one obtains that, in the definition of \mathcal{Q}_n , it can be assumed that the indices *i* are restricted to the set $\{[\rho], \ldots, \beta_n\}$.

Let K be such that the support of f is included in $[0, K]^2$, if $u_n = \log \log n$, for $i \ge \lceil \rho \rceil$,

$$\mathbb{E}\left(f(i/n^{1/(\rho+2)}, nP_i)\mathbb{1}_{\{t_{\lfloor\rho\rfloor} \le u_n\}}\right) \le \|f\|_{\infty} \mathbb{P}\left(t_{\lfloor\rho\rfloor} \le u_n, nP_i \le K\right),$$

since $P_i = e^{-\rho t_{\lfloor \rho \rfloor}} e^{-\rho (t_{i-1} - t_{\lfloor \rho \rfloor})} (1 - e^{-\rho E_i/i}),$

$$\begin{split} \mathbb{E}\left(f(i/n^{1/(\rho+2)}, nP_i)\mathbb{1}_{\{t_{\lfloor\rho\rfloor} \le u_n\}}\right) \\ & \leq \|f\|_{\infty} \mathbb{P}\left[\left(\frac{1-e^{-\rho E_i/i}}{\rho/i}\right) \le \frac{i}{\rho} K e^{\rho u_n} e^{\rho(t_{i-1}-t_{\lfloor\rho\rfloor})}/n\right]. \end{split}$$

By using the elementary inequality, if E_1 is exponentially distributed with mean 1,

(8)
$$\mathbb{P}\left(\frac{1}{y}\left(1-e^{-yE_1}\right) \le x\right) \le e\left(1-e^{-x}\right), \quad y \le 1, x \ge 0,$$

one gets that, for $i > \rho$,

$$\begin{split} \mathbb{E}\left(f(i/n^{1/(\rho+2)}, nP_i)\mathbb{1}_{\{t_{\lfloor\rho\rfloor} \le u_n\}}\right) \\ & \le e\|f\|_{\infty} \mathbb{E}\left(1 - \exp\left[-\frac{i}{n\rho}Ke^{\rho u_n}e^{\rho(t_{i-1}-t_{\lfloor\rho\rfloor})}\right]\right) \\ & \le eK\|f\|_{\infty} \frac{ie^{\rho u_n}}{n\rho} \mathbb{E}\left(e^{\rho(t_{i-1}-t_{\lfloor\rho\rfloor})}\right) \\ & = eK\|f\|_{\infty} \frac{ie^{\rho u_n}}{n\rho}e^{\rho\sum_{k=\lceil\rho\rceil}^{i-1}1/k}e^{\sum_{k=\lceil\rho\rceil}^{i-1}-\log(1-\rho/k)-\rho/k}. \end{split}$$

Thus, there exists some finite constant C such that, for $i > \rho$,

$$\mathbb{E}\left(f(i/n^{1/(\rho+2)}, nP_i)\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le u_n\}}\right) \le C \frac{i^{\rho+1}e^{\rho u_n}}{n} = C \frac{i^{\rho+1}(\log n)^{\rho}}{n},$$

consequently,

$$\mathbb{E}\left(\mathcal{Q}_n(f)\mathbb{1}_{\{t_{\lfloor\rho\rfloor} \le u_n\}}\right) \le C\frac{\beta_n^{\rho+2}(\log n)^{\rho}}{n} \le C\frac{1}{(\log n)^2}$$

This relation and the inequality

$$\mathbb{E}\left(1-e^{-\mathcal{Q}_n(f)}\right) \le \mathbb{P}(t_{\lfloor \rho \rfloor} > u_n) + \mathbb{E}\left(\mathcal{Q}_n(f)\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le u_n\}}\right)$$

give the desired result.

Step 3. The proof of the theorem can be now completed. By Equation (3), for $i \geq 1$, $P_i = W_i^{\rho} Z_i / i^{\rho+1}$, by using Step 2 and the same techniques as in Step 1 together with the fact that, for $\eta > 0$, the probability of the event

$$\left\{\sup\left(\left|W_{i}^{\rho}/W_{\beta_{n}}^{\rho}-1\right|:i\geq\beta_{n}\right)\geq\eta\right\}$$

converges to 0 as n gets large, it is not difficult to show that the sequences of point processes

$$\left\{ \left(\frac{i}{n^{1/(\rho+2)}}, \frac{n}{i^{\rho+1}} W_i^{\rho} Z_i \right), i \ge 1 \right\} \text{ and } \left\{ \left(\frac{i}{n^{1/(\rho+2)}}, \frac{n}{i^{\rho+1}} W_{\beta_n}^{\rho} Z_i \right), i \ge \beta_n \right\}$$

have the same limit in distribution. Because W_{β_n} is independent of $(Z_i, i \ge \beta_n)$, the last point process has the same distribution as \mathcal{P}_n^2 (up to the first β_n terms which are negligible similarly as in Step 2). By Step 1, the convergence in distribution is therefore proved.

The following proposition strengthens the statement of Proposition 3.1, it will be used to prove the main asymptotic result on the indices of empty bins.

PROPOSITION 3.3. If $f : \mathbb{R}^2_+ \to \mathbb{R}_+$ is a continuous function such that

- (1) there exists K such that f(x, y) = 0 for any $x \ge K$ and $y \in \mathbb{R}_+$,
- (2) for all $x \in \mathbb{R}_+$, the function $y \mapsto f(x, y)$ is differentiable and

$$y \mapsto y \left\| \frac{\partial f}{\partial y} \right\|_{y} \stackrel{def.}{=} y \sup_{x \in \mathbb{R}_{+}} \left| \frac{\partial f}{\partial y} \right| (x, y)$$

is integrable on \mathbb{R}_+ ,

then Convergence (7) also holds for f.

PROOF. For $M, L \ge 0$ and $i, n \in \mathbb{N}$, one has

$$\mathbb{E}\left(f\left(\frac{i}{n^{\rho+2}}, nP_i\right)\mathbb{1}_{\{nP_i \ge M, t_{\lfloor\rho\rfloor} \le L\}}\right)$$
$$= -\int_0^{+\infty} \frac{\partial f}{\partial y}\left(\frac{i}{n^{\rho+2}}, y\right)\mathbb{P}(M \le nP_i \le y, t_{\lfloor\rho\rfloor} \le L)\,dy.$$

By using similar arguments as in the end of the proof of the above theorem, one gets

$$\mathbb{E}\left(f\left(\frac{i}{n^{\rho+2}}, nP_{i}\right)\mathbb{1}_{\{nP_{i}\geq M, t_{\lfloor\rho\rfloor}\leq L\}}\right)$$

$$\leq e\int_{M}^{+\infty} \left\|\frac{\partial f}{\partial y}\right\|_{y} \mathbb{E}\left(1-\exp\left[-\frac{i}{n\rho}ye^{\rho L}e^{\rho(t_{i-1}-t_{\lfloor\rho\rfloor})}\right]\right) dy$$

$$\leq \frac{ie^{\rho L}}{n\rho}e\mathbb{E}\left(e^{\rho(t_{i-1}-t_{\lfloor\rho\rfloor})}\right)\int_{M}^{+\infty}y\left\|\frac{\partial f}{\partial y}\right\|_{y} dy$$

$$\leq C\frac{i^{\rho+1}e^{\rho L}}{n}\int_{M}^{+\infty}y\left\|\frac{\partial f}{\partial y}\right\|_{y} dy,$$

for some fixed constant C. Define $k_n = \lfloor Kn^{1/(\rho+2)} \rfloor$, by summing up these terms, this gives the relation

(9)
$$\mathbb{E}\left(\sum_{i\geq 1} f\left(\frac{i}{n^{\rho+2}}, nP_i\right) \mathbb{1}_{\{M\leq nP_i, t_{\lfloor\rho\rfloor}\leq L\}}\right)$$
$$\leq C \frac{k_n^{\rho+2} e^{\rho L}}{n} \int_M^{+\infty} y \left\|\frac{\partial f}{\partial y}\right\|_y dy \leq C K^{\rho+2} e^{\rho L} \int_M^{+\infty} y \left\|\frac{\partial f}{\partial y}\right\|_y dy.$$

Define $f_0(x, y) = f(x, y) \mathbb{1}_{\{y \le M\}}$, by using a convolution kernel on the variable y, there exist sequences (g_p^+) and (g_p^-) in $C_c^+(\mathbb{R}_+)$ converging pointwisely to f_0 for all $y \ne M$ such that $g_p^- \le f_0 \le g_p^+$. See Rudin [**Rud87**] for example. Proposition 3.1 gives that

$$\mathbb{E}(\exp(-\mathcal{P}(g_p^+))) \leq \liminf_{n \to +\infty} \mathbb{E}(\exp(-\mathcal{P}_n(f_0)))$$
$$\leq \limsup_{n \to +\infty} \mathbb{E}(\exp(-\mathcal{P}_n(f_0))) \leq \mathbb{E}(\exp(-\mathcal{P}(g_p^-))),$$

and Expression (6) shows that, as p goes to infinity, the left and right hand side terms of this relation converge to the Laplace transform of \mathcal{P} at f_0 . Therefore, Convergence (7) holds at f_0 . Since

$$\begin{aligned} 0 &\leq \mathbb{E}\left(e^{-\mathcal{P}_{n}(f)}\right) - \mathbb{E}\left(e^{-\mathcal{P}_{n}(f_{0})}\right) \\ &\leq P(t_{\lfloor\rho\rfloor} \geq L) + \mathbb{E}\left[\left(1 - e^{-(\mathcal{P}_{n}(f) - \mathcal{P}_{n}(f_{0}))}\right)\mathbbm{1}_{\{t_{\lfloor\rho\rfloor} \leq L\}}\right] \\ &\leq P(t_{\lfloor\rho\rfloor} \geq L) + \mathbb{E}\left[\left(\mathcal{P}_{n}(f) - \mathcal{P}_{n}(f_{0})\right)\mathbbm{1}_{\{t_{\lfloor\rho\rfloor} \leq L\}}\right], \end{aligned}$$

and the last term being the left hand side of Relation (9), one can choose L and M sufficiently large so that this difference is arbitrarily small. The proposition is proved.

4. Asymptotic Behavior of the Indices of the First Empty Bins

It is assumed that a large number n of balls are thrown in the bins according to the probability distribution (P_i) defined by Equation (3). The purpose of this

section is to establish limit theorems to describe the limiting distribution of the set of indices of bins having a fixed number of balls.

THEOREM 4.1. The point process of rescaled indices of empty bins associated to the probability vector (P_i) when n balls have been used

$$\mathcal{N}_n = \left\{ \frac{i}{n^{1/(\rho+2)}} : i \ge 1, \text{ the ith bin is empty} \right\}$$

converges in distribution as n goes to infinity to a point process \mathcal{N}_{∞} whose distribution is given by

$$\mathbb{E}\left(e^{-\mathcal{N}_{\infty}(g)}\right) = \mathbb{E}\left[\exp\left(-\frac{W_{\infty}^{-\rho}}{\rho}\int_{\mathbb{R}_{+}}\left(1-e^{-g(x)}\right)x^{\rho+1}\,dx\right)\right],$$

for $g \in C_c^+(\mathbb{R}_+)$. Equivalently (\mathcal{N}_n) converges in distribution to the point process

$$\left(W_{\infty}^{\rho/(\rho+2)} t_i^{1/(\rho+2)}\right)$$

where (t_i) is a standard Poisson process with parameter $[\rho(\rho+2)]^{-1/(\rho+2)}$.

It can also be shown that the same result holds when the indices of bins containing k balls are considered. If $\mathcal{N}_{k,n}$ is the corresponding point process, the limiting point process does not in fact depend on k and, moreover, the sequence $(\mathcal{N}_{k,n}, k \ge 0)$ converges in distribution to $(\mathcal{N}_{k,\infty}, k \ge 0)$ and, conditionally on W_{∞} , the random variables $\mathcal{N}_{k,\infty}, k \ge 0$ are independent with the same distribution.

PROOF. The proof uses the poissonization technique: A closely related model is first analyzed when U_n balls are used, U_n being an independent Poisson random variable with mean n, \mathcal{N}_n^0 denotes the corresponding point process. For this model, conditionally on the sequence (P_i) , the number of balls in the bins are independent Poisson random variables with respective parameters (nP_i) . In a first step, the convergence in distribution of the sequence (\mathcal{N}_n^0) of point processes is investigated. Let $g \in C_c^+(\mathbb{R}_+)$,

$$\mathbb{E}\left(e^{-\mathcal{N}_{n}^{0}(g)}\right) = \mathbb{E}\left(\exp\left[\sum_{i=1}^{+\infty}\log\left[1 - e^{-nP_{i}}\left(1 - e^{-g(i/n^{1/(\rho+2)})}\right)\right]\right]\right),$$

if one defines $f(x, y) = -\log \left[1 - e^{-y} \left(1 - e^{-g(x)}\right)\right]$, then

 $\mathbb{E}\left(\exp\left[-\mathcal{N}_{n}^{0}(g)\right]\right) = \mathbb{E}\left(\exp\left[-\mathcal{P}_{n}(f)\right]\right),$

where \mathcal{P}_n is the point process defined in Theorem 3.1. By using Proposition 3.3, one gets the relation

$$\lim_{n \to +\infty} \mathbb{E}\left(e^{-\mathcal{N}_{n}^{0}(g)}\right) = \mathbb{E}\left[\exp\left(-\frac{W_{\infty}^{-\rho}}{\rho}\int_{\mathbb{R}^{2}_{+}}\left(1-e^{-f(x,y)}\right)x^{\rho+1}\,dx\,dy\right)\right]$$
$$= \mathbb{E}\left[\exp\left(-\frac{W_{\infty}^{-\rho}}{\rho}\int_{\mathbb{R}^{+}}\left(1-e^{-g(x)}\right)x^{\rho+1}\,dx\right)\right].$$

Consequently, the sequence (\mathcal{N}_n^0) converges in distribution to \mathcal{N}_∞ . For $0 < \alpha < 1$, it is not difficult to check that the same convergence result holds when $U_{n+n^{\alpha}}$ balls are used: If \mathcal{N}_n^1 denotes the associated point process, then (\mathcal{N}_n^1) converges also in

distribution to \mathcal{N}_{∞} . For x > 0, the monotonicity property $\mathcal{N}_a([0, x]) \leq \mathcal{N}_b([0, x])$ for $b \leq a$ gives the relation

$$\mathbb{P}(\mathcal{N}_n^1([0,x]) \le k) \le \mathbb{P}\left(\mathcal{N}_n([0,x]) \le k\right) + \mathbb{P}(U_{n+n^{\alpha}} \le n).$$

The central limit theorem for Poisson processes shows that for $\alpha \in (1/2, 1)$, the quantity $\mathbb{P}(U_{n+n^{\alpha}} \leq n)$ converges to 0 as n gets large, therefore if $k \geq 0$,

$$\mathbb{P}(\mathcal{N}_{\infty}([0,x]) \le k) = \lim_{n \to +\infty} \mathbb{P}(\mathcal{N}_{n}^{1}([0,x]) \le k) \le \liminf_{n \to +\infty} \mathbb{P}\left(\mathcal{N}_{n}([0,x]) \le k\right).$$

By using a similar argument with the lim sup, one gets that the sequence $(\mathcal{N}_n([0, x]))$ converges in distribution to $\mathcal{N}_{\infty}([0, x])$. With the same coupling argument, one gets that for $x_1 \leq x_2 \leq \cdots \leq x_p \in \mathbb{R}_+$ and $(k_i, 1 \leq i \leq p) \in \mathbb{N}^p$,

$$\lim_{n \to +\infty} \mathbb{P}(\mathcal{N}_n([0, x_1]) \le k_1, \mathcal{N}_n([0, x_2]) \le k_2, \dots, \mathcal{N}_n([0, x_p]) \le k_p)$$
$$= \mathbb{P}(\mathcal{N}_\infty([0, x_1]) \le k_1, \mathcal{N}_\infty([0, x_2]) \le k_2, \dots, \mathcal{N}_\infty([0, x_p]) \le k_p).$$

and therefore the convergence in distribution of (\mathcal{N}_n) to \mathcal{N}_{∞} . The proposition is proved.

COROLLARY 4.1. If ν_n is the index of the first empty bin when n balls are thrown, then

$$\lim_{n \to +\infty} \mathbb{P}\left(\frac{\nu_n}{n^{1/(\rho+2)}} \ge x\right) = \mathbb{E}\left(\exp\left(-\frac{x^{\rho+2}W_{\infty}^{-\rho}}{\rho(\rho+2)}\right)\right), \quad x \ge 0.$$

Comparison with Deterministic Power Law Decay. For $\delta > 1$, one considers the bins and balls problem with the probability vector $Q = (Q_i, i \ge 1) = (\alpha/i^{\delta})$. Note that for the problems analyzed in this paper, only the asymptotic behavior of the sequence (Q_i) matters. The equivalent of Theorem 4.1 can be obtained from Theorem III.4.1.

PROPOSITION 4.1. As n goes to infinity, the point process

$$\left\{\frac{i(\log n)^{1/\delta+1}}{(\alpha\delta n)^{1/\delta}} - \log n - \frac{1+\delta}{\delta}\log\log n: \text{ the ith bin is empty}\right\}$$

converges in distribution to a Poisson point process with the intensity measure $(\alpha\delta)^{1/\delta}e^x dx$ on \mathbb{R} .

The probability vector considered in the above theorem has an asymptotic expression of the form $(P_i) = (W_{\infty}^{\rho}F_i/i^{\rho+1})$. In this case, empty bins show up for indices of the order of $n^{1/(\rho+2)}$, i.e., much earlier than for the deterministic case where the exponent of n is $1/\delta = 1/(\rho + 1)$ (if one ignores the log). This can be explained simply by the fact that some of the i.i.d. exponential random variables (F_i) can be very small thereby creating an additional possibility of having empty bins.

In this picture, the variable W_{∞} does not seem to have an influence on the qualitative behavior of these occupancy schemes other than creating some dependency structure for the vector (P_i) . The next section shows that this variable has nevertheless an important role if one looks at the averages of the number of empty bins.

5. Rare Events

From Equation (7) of Theorem 3.1, for x > 0, the limiting number (in distribution) of empty bins whose index is less than $xn^{1/(\rho+2)}$ has an average value given by

$$\frac{x^{\rho+2}}{\rho(\rho+2)}\mathbb{E}\left(W_{\infty}^{-\rho}\right) = \frac{x^{\rho+2}}{\rho(\rho+2)}\int_{0}^{+\infty}\frac{1}{u^{\rho}}e^{-u}\,du,$$

by Equation (2) and since $W_{\infty} = \exp(-M_{\infty})$. This quantity is infinite when $\rho \geq 1$. The goal of this section is to investigate this phenomenon which has a significant impact on the peer-to-peer system of Chapter III, at the origin of this model. For this purpose a family of rescaled point processes is introduced.

DEFINITION 3. If $\phi : \mathbb{N} \to \mathbb{R}_+$ is a non-decreasing function, for $n \ge 1$, \mathcal{N}_n^{ϕ} denotes the point process defined by

$$\mathcal{N}_n^{\phi} = \left\{ \frac{i}{\phi(n)} : i \ge 1, \text{ the } i \text{th bin is empty} \right\}.$$

For $i \ge 1$, recall that, by Equation (3), the probability P_i of throwing a ball in the *i*th bin is, with the notations of Section 2,

$$P_i = \frac{1}{i^{\rho+1}} W_i^{\rho} Z_i$$
, with $Z_i = i \left(1 - e^{-\rho E_i/i} \right)$ and $W_i = e^{-M_{i-1}}$,

where (E_i) are i.i.d. exponential random variables with parameter 1, and define its asymptotic representation as

$$\widetilde{P}_i \stackrel{\text{def.}}{=} \frac{1}{i^{\rho+1}} W^{\rho}_{\infty} F_i$$
, with $W_{\infty} = e^{-M_{\infty}}$,

where W_i and W_{∞} are independent of (F_i) , an i.i.d. sequence of exponential random variables with mean ρ .

The following proposition shows that, as in the proof of Theorem 4.1, for some asymptotics for (\mathcal{N}_n^{ϕ}) , it is enough to analyze the asymptotic behaviour of a functional of the sequence (\widetilde{P}_i) .

PROPOSITION 5.1. For x > 0 and a non-decreasing function ϕ , then

$$\lim_{n \to +\infty} \mathbb{E}\left(\mathcal{N}_{n}^{\phi}([0,x])\right) / \sum_{i=1}^{\lfloor x\phi(n) \rfloor} \mathbb{E}\left(e^{-n\tilde{P}_{i}}\right) = 1,$$

provided that

$$\lim_{n \to +\infty} \frac{\phi(n)}{n^{\alpha_0}} = 0 \text{ and } \lim_{n \to +\infty} \sum_{i=1}^{\lfloor x\phi(n) \rfloor} \mathbb{E}\left(e^{-n\tilde{P}_i}\right) = +\infty.$$

for some $\alpha_0 < 1/(\rho + 1)$.

PROOF. See Appendix A.

For $i > \lfloor \rho \rfloor$, the quantity P_i will be written as $P_i = \exp(-\rho t_{\lfloor \rho \rfloor}) D_i Z_i / i^{\rho+1}$ with

$$D_{i} \stackrel{\text{def.}}{=} \exp\left(-\rho\left[M_{i-1} - M_{\lfloor \rho \rfloor} - \log\left\lfloor \rho + 1 \rfloor\right]\right)$$

The sequence (D_i) converges almost surely to a finite limit D_{ρ} given by

(10)
$$D_{\rho} \stackrel{\text{def.}}{=} \exp\left(-\rho \left[M_{\infty} - M_{\lfloor \rho \rfloor} - \log \lfloor \rho + 1 \rfloor\right]\right),$$

and, since $\exp(\rho E_i/i)$ is integrable for $i > \rho$, a similar result holds for the expected values

$$\lim_{i \to +\infty} \mathbb{E}(1/D_i) = \mathbb{E}(1/D_\rho) < +\infty.$$

With this definition, the sequence (\widetilde{P}_i) can be represented as

$$\left(\widetilde{P}_{i}\right) = \left(\frac{e^{-\rho t_{\lfloor \rho \rfloor}}}{i^{\rho+1}}D_{\rho}F_{i}\right).$$

Some Preliminary Estimates. For the moment, $k \in \mathbb{N}$ is fixed, if $n \ge 1$, $i > \rho$, then

$$\mathbb{E}\left(e^{-n\tilde{P}_{i}}\right) = \mathbb{E}\left[\exp\left(-n\rho D_{\rho}e^{-\rho t_{\lfloor\rho\rfloor}}F_{1}/i^{\rho+1}\right)\right]$$
$$= \mathbb{E}\left(\frac{i^{\rho+1}/n}{i^{\rho+1}/n + e^{-\rho t_{\lfloor\rho\rfloor}}\rho D_{\rho}}\right),$$

by summing up these terms, if $\varepsilon_{k,n} \stackrel{\text{def.}}{=} k/n^{1/(\rho+1)}$, one gets that

$$\sum_{i=\lfloor\rho\rfloor+1}^{k} \mathbb{E}\left(e^{-n\widetilde{P}_{i}}\right) = n^{1/(\rho+1)} \int_{0}^{\varepsilon_{k,n}} \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1} + e^{-\rho t_{\lfloor\rho\rfloor}}\rho D_{\rho}}\right) dv + O\left(\varepsilon_{k,n}\right),$$

which gives the relation

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\tilde{P}_{i}}\right) = n^{1/(\rho+1)} \varepsilon_{k,n}^{\rho+2} \int_{0}^{1} \mathbb{E}\left(\frac{v^{\rho+1}}{\varepsilon_{k,n}^{\rho+1} v^{\rho+1} + e^{-\rho t_{\lfloor \rho \rfloor}} \rho D_{\rho}}\right) dv + O\left(\varepsilon_{k,n}\right),$$

with a change of variables. By using Equation (1) and again a change of variables, one obtains the relation

(11)
$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\tilde{P}_{i}}\right) = \frac{\lfloor\rho\rfloor}{\rho} n^{1/(\rho+1)} \varepsilon_{k,n}^{(2\rho+1)/\rho} \times \int_{0}^{1/\varepsilon^{\rho+1}} u^{1/\rho-1} (1-\varepsilon_{k,n}^{\frac{\rho+1}{\rho}} u^{1/\rho})^{\lfloor\rho\rfloor-1} du \int_{0}^{1} \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1}+u\rho D_{\rho}}\right) dv + O\left(\varepsilon_{k,n}\right).$$

This quantity is now analyzed according to the values of ρ .

Case $\rho > 1$.

If $k_n = \lfloor xn^{\alpha} \rfloor$ with $1/(2\rho+1) \leq \alpha < 1/(\rho+1)$, then $\varepsilon_{k_n,n} \sim xn^{(\alpha(\rho+1)-1)/(\rho+1)}$ and, by Relation (11),

$$\lim_{n \to +\infty} \frac{1}{n^{((2\rho+1)\alpha-1)/\rho}} \sum_{i=1}^{k_n} \mathbb{E}\left(e^{-n\tilde{P}_i}\right)$$
$$= x^{(2\rho+1)/\rho} \frac{\lfloor\rho\rfloor}{\rho} \int_0^{+\infty} u^{1/\rho-1} du \int_0^1 \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1} + u\rho D_\rho}\right) dv.$$

Case $\rho = 1$. Equation (11) is for this case

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\tilde{P}_{i}}\right) = \sqrt{n}\varepsilon_{k,n}^{3} \int_{0}^{1/\varepsilon_{k,n}^{2}} du \int_{0}^{1} \mathbb{E}\left(\frac{v^{2}}{v^{2} + uD_{1}}\right) dv + O\left(\varepsilon_{k,n}\right).$$

If $k_n = |xn^{1/3}/\log^\beta n|$ with $\beta \in \mathbb{R}$, then $\varepsilon_{k_n,n} \sim x/(n^{1/6}(\log n)^\beta)$ and for $\beta \leq 1/3$,

$$\lim_{n \to +\infty} \frac{1}{(\log n)^{(1-3\beta)}} \sum_{i=1}^{k_n} \mathbb{E}\left(e^{-n\tilde{P}_i}\right) = \frac{1}{9} x^3 \mathbb{E}\left(\frac{1}{D_1}\right).$$

The following proposition has therefore been proved.

PROPOSITION 5.2 (Average of the Number of Empty Bins). For α , $\beta > 0$, for $n \in \mathbb{N}$, denote by $p_{\alpha,\beta}(n) = n^{\alpha} (\log n)^{-\beta}$, and by convention $p_{\alpha} = p_{\alpha,0}$.

(1) If
$$\rho > 1$$
 and $1/(2\rho + 1) \le \alpha < 1/(\rho + 1)$.

$$\lim_{n \to +\infty} \frac{1}{n^{(\alpha(2\rho+1)-1)/\rho}} \mathbb{E} \left(\mathcal{N}_n^{p_\alpha}([0,x]) \right) \\ = x^{(2\rho+1)/\rho} \frac{\lfloor \rho \rfloor}{\rho} \int_0^{+\infty} u^{1/\rho-1} du \int_0^1 \mathbb{E} \left(\frac{v^{\rho+1}}{v^{\rho+1} + u\rho D_\rho} \right) dv.$$
(2) If $\rho = 1$ and $\beta \le 1/3$,

$$\lim_{n \to +\infty} \frac{1}{(\log n)^{(1-3\beta)}} \mathbb{E} \left(\mathcal{N}^{p_{1/3,\beta}}([0,x]) \right) = \frac{1}{9} x^3 \mathbb{E} \left(\frac{1}{D_1} \right).$$

A Double Threshold. For the convergence in distribution of the sequence of point processes (\mathcal{N}_n^{ϕ}) , Theorem 4.1 has shown that the correct scaling ϕ for the order of magnitude of the indices of the first empty bins is given by $\phi(n) = n^{1/(\rho+2)}, n \ge 1$. For the average number of points in a finite interval, the above proposition states that, for $\rho > 1$, the correct scaling is in fact $\phi(n) = n^{1/(2\rho+1)} \ll n^{1/(\rho+2)}$.

For $\alpha > 0$, with the notations of the above proposition, one concludes that under the condition $\rho > 1$ and for $1/(2\rho + 1) < \alpha < 1/(\rho + 2)$, the following limit results hold

$$\mathcal{N}_n^{p_\alpha} \stackrel{\text{dist.}}{\to} 0 \text{ and } \lim_{n \to +\infty} \mathbb{E} \left(\mathcal{N}_n^{p_\alpha}[0, x] \right) = +\infty, \quad \forall x > 0.$$

This suggests that, in this case, with a high probability, all the bins with index less than $n^{1/(\rho+2)}$ have a large number of balls. But also that there exists some rare event for which a very large number of empty bins with indices of an order slightly greater than $n^{1/(2\rho+1)}$ are created. The following proposition shows that the total size of the first $|\rho|$ bins is the key variable to explain this phenomenon. It should be of the order of $\log n$ in order to have sufficiently many empty bins in the appropriate region.

PROPOSITION 5.3. For $\rho > 1$ and if $p_{\alpha}(n) = n^{\alpha}$, for $\alpha \in [1/(2\rho+1), 1/(\rho+2))$ and

$$\delta_0(\alpha) \stackrel{\text{def.}}{=} \frac{1 - \alpha(\rho + 2)}{\rho - 1} \quad and \quad \delta_1(\alpha) \stackrel{\text{def.}}{=} \frac{1 - \alpha(\rho + 1)}{\rho},$$

then, for $a \in \mathbb{R}$ and x > 0,

(1) If $\delta < \delta_0(\alpha)$, then

$$\lim_{n \to +\infty} \mathbb{E}\left(\mathcal{N}_n^{p_{\alpha}}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le \delta \log n\}}\right) = 0.$$

(2) If
$$\delta \in [\delta_0(\alpha), \delta_1(\alpha)]$$
, then

$$\lim_{n \to +\infty} \frac{\mathbb{E}\left(\mathcal{N}_n^{p_\alpha}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le \delta \log n+a\}}\right)}{n^{(\rho+2)\alpha+\delta(\rho-1)-1}} = \frac{x^{\rho+2}}{(\rho+2)} \frac{\lfloor \rho \rfloor}{(\rho-1)} \mathbb{E}\left(\frac{1}{\rho D_\rho}\right) e^{(\rho-1)a}$$

(3) If
$$\delta \ge \delta_1(\alpha)$$
,

$$\lim_{n \to +\infty} \frac{\mathbb{E}\left(\mathcal{N}_n^{p_\alpha}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le \delta \log n+a\}}\right)}{n^{((2\rho+1)\alpha-1)/\rho}}$$

$$= x^{(2\rho+1)/\rho} \frac{\lfloor \rho \rfloor}{\rho} \int_{e^{-\rho\alpha}\mathbb{1}_{\{\delta = \delta_1(\alpha)\}}}^{+\infty} u^{1/\rho-1} du \int_0^1 \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1} + u\rho D_\rho}\right) dv,$$

where D_{ρ} is the random variable defined by Equation (10).

PROOF. To begin with, it is assumed that $\delta \in [\delta_0(\alpha), \delta_1(\alpha))$. If $k \ge 1, b > 0$, $\varepsilon_{k,n} = k/n^{1/(\rho+1)}$, $k = \lfloor xn^{\alpha} \rfloor$ and $b = \delta \log n + a$, in the same way as for Equation (11), one gets

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\widetilde{P}_{i}}\mathbb{1}_{\left\{t_{\lfloor\rho\rfloor} \leq b\right\}}\right) = \frac{\lfloor\rho\rfloor}{\rho} n^{1/(\rho+1)} \varepsilon_{k,n}^{(2\rho+1)/\rho} \times \int_{e^{-\rho b}/\varepsilon_{k,n}^{\rho+1}}^{1/\varepsilon_{k,n}^{\rho+1}} u^{1/\rho-1} (1 - \varepsilon_{k,n}^{\frac{\rho+1}{\rho}} u^{1/\rho})^{\lfloor\rho\rfloor - 1} du \int_{0}^{1} \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1} + u\rho D_{\rho}}\right) dv + O(\varepsilon_{k,n})$$
(12)

Note that

$$e^{-\rho b}/\varepsilon_{k,n}^{\rho+1} \sim n^{1-\rho\delta-\alpha(\rho+1)}e^{-\rho a} \nearrow +\infty$$

hence the range of the first integral goes to infinity as n gets large. Since

$$\int_0^1 \mathbb{E}\left(\frac{uv^{\rho+1}}{v^{\rho+1}+u\rho D_\rho} - \frac{v^{\rho+1}}{\rho D_\rho}\right) dv = \int_0^1 \mathbb{E}\left(\frac{v^{2(\rho+1)}}{\rho(v^{\rho+1}+u\rho D_\rho)D_\rho}\right) dv,$$

by Lebesgue's Theorem, this integral is arbitrarily small as \boldsymbol{u} gets large, this implies the equivalence

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\widetilde{P}_{i}}\mathbb{1}_{\{t_{\lfloor\rho\rfloor} \le b\}}\right) \sim \frac{\lfloor\rho\rfloor}{\rho(\rho+2)} \mathbb{E}\left(\frac{1}{\rho D_{\rho}}\right) n^{1/(\rho+1)} \varepsilon_{k,n}^{(2\rho+1)/\rho} \int_{e^{-\rho b}/\varepsilon_{k,n}^{\rho+1}}^{1/\varepsilon_{k,n}^{\rho+1}} u^{1/\rho-2} du$$

If C is the multiplicative constant of the right hand side of the above relation, then

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\widetilde{P}_{i}}\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \leq b\}}\right) \sim \frac{C\rho}{\rho-1} \frac{k^{\rho+2}}{n} \left(e^{b(\rho-1)} - 1\right),$$

this gives the equivalence

$$\sum_{i=1}^{k} \mathbb{E}\left(e^{-n\tilde{P}_{i}}\mathbbm{1}_{\{t_{\lfloor\rho\rfloor} \leq b\}}\right) \sim x^{\rho+2} \frac{C\rho}{\rho-1} e^{a(\rho-1)} n^{(\rho+2)\alpha+\delta(\rho-1)-1} d\rho^{-1} d\rho^{-1}$$

The proof of this case is completed.

The case $\delta \geq \delta_1(\alpha)$ uses Equation (12). The term $e^{-\rho b}/\varepsilon_{k,n}^{\rho+1}$ converges to $e^{-\rho a}$ if $\delta = \delta_1(\alpha)$ and 0 otherwise. This gives directly the desired convergence.

Finally, if $\delta < \delta_0(\alpha)$, for any $a \in \mathbb{R}$, there exists n_0 so that if $n \ge n_0$, then $\delta \log n \le \delta_0(\alpha) \log n + a$, in particular

$$\mathbb{E}\left(\mathcal{N}_{n}^{p_{\alpha}}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \leq \delta \log n\}}\right) \leq \mathbb{E}\left(\mathcal{N}_{n}^{p_{\alpha}}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \leq \delta_{0}(\alpha) \log n+a\}}\right)$$

hence

$$\limsup_{n \to +\infty} \mathbb{E}\left(\mathcal{N}_n^{p_{\alpha}}([0,x])\mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le \delta \log n\}}\right) \le \frac{x^{\rho+2}}{(\rho+2)} \frac{\lfloor \rho \rfloor}{(\rho-1)} \mathbb{E}\left(\frac{1}{\rho D_{\rho}}\right) e^{(\rho-1)a}.$$

One concludes by letting a go to $-\infty$.

As a consequence of the above proposition, for $\alpha \in [1/(2\rho + 1), 1/(\rho + 2))$, the average of the variable $\mathcal{N}_n^{p_\alpha}([0, x])$ converges to infinity only when the total size $t_{\lfloor \rho \rfloor}$ of the first $\lfloor \rho \rfloor$ bins is of the order $\delta \log n$ for a sufficiently large δ . The following corollary gives a more precise formulation.

COROLLARY 5.1. For $\rho > 1$ and if $p_{\alpha}(n) = n^{\alpha}$, for $\alpha \in [1/(2\rho + 1), 1/(\rho + 2))$ $\delta_{\alpha}(\alpha) = (1 - \alpha(\alpha + 1))/\alpha$

$$\delta_1(\alpha) = (1 - \alpha(\rho + 1))/\rho$$

then, for a, b > 0,

$$\lim_{n \to +\infty} \frac{\mathbb{E}\left(\mathcal{N}_n^{p_\alpha}([0,x])\mathbb{1}_{\{\delta_1(\alpha)\log n - a \le t_{\lfloor\rho\rfloor} \le \delta_1(\alpha)\log n + b\}}\right)}{\mathbb{E}\left(\mathcal{N}_n^{p_\alpha}([0,x])\right)} = \psi(-a,b)$$

where, for $y,\,z\in\mathbb{R},\,\psi(y,z)=\phi(y,z)/\phi(-\infty,+\infty)$ and

$$\phi(y,z) = x^{(2\rho+1)/\rho} \frac{\lfloor \rho \rfloor}{\rho} \int_{[e^{-\rho z}, e^{-\rho y}]} u^{1/\rho-1} du \int_0^1 \mathbb{E}\left(\frac{v^{\rho+1}}{v^{\rho+1} + u\rho D_\rho}\right) dv.$$

A rough (non-rigorous) interpretation of this result could be as follows: on the event where "most" (i.e., for the averages) of empty bins are created in the interval $[0, xn^{\alpha}]$, the random variable $t_{\lfloor \rho \rfloor} - \delta_1(\alpha) \log n$ converges in distribution to some random variable X on \mathbb{R} , such that $\mathbb{P}(X \leq a) = \psi(-\infty, a)$.

The following analogous result is proved in a similar way for the critical case $\rho = 1$.

PROPOSITION 5.4. For $\rho = 1$ and with the notations of the above proposition then, for $0 < \beta < 1/3$, x > 0, and for $0 \le a \le 1/3$,

$$\lim_{n \to +\infty} \frac{1}{(\log n)^{1-3\beta}} \mathbb{E}\left(\mathcal{N}_n^{p_{1/3,\beta}}([0,x]) \mathbb{1}_{\{t_{\lfloor \rho \rfloor} \le a \log n\}}\right) = \frac{a}{3} x^3 \mathbb{E}\left(\frac{1}{D_1}\right),$$

and for a > 1/3,

$$\lim_{n \to +\infty} \frac{1}{(\log n)^{1-3\beta}} \mathbb{E}\left(\mathcal{N}_n^{p_{1/3,\beta}}([0,x])\mathbbm{1}_{\{t_{\lfloor \rho \rfloor} \le a \log n\}}\right) = \frac{1}{9} x^3 \mathbb{E}\left(\frac{1}{D_1}\right),$$

where D_1 is the random variable defined by Equation (10).

6. Generalizations

The problem analyzed in the present chapter can be generalized towards two directions. On the one hand, the sequence (t_n) can stem from a general branching process instead of the particular Yule one; on the other hand, the locations of balls can have a general distribution. This section discusses these possible extensions.

Exponential Balls and General Branching Process. Let (t_n) be the birth instants of a general supercritical branching process (Z(t)). See Kingman **[Kin75]** and Nerman **[Ner81]** for example. Let α be the Malthusian parameter, and W the almost sure limit of $(e^{-\alpha t}Z(t))$. Under reasonable technical assumptions, Härnqvist **[H81]** has shown the following result:

THEOREM 6.1. Define the point process Ψ_t^* by

$$\Psi_t = \sum_{k \ge 1} \mathbb{1}_{\{t \le t_k\}} \delta_{t_k e^{\alpha t}},$$

as t gets large, Ψ_t converges in distribution to a mixed Poisson process whose parameter is distributed as γW for some constant $\gamma > 0$.

From this result, it is possible to prove that the process $(n(t_{n+k} - t_n), k \ge 1)$ converges in distribution, as n goes to infinity, to a Poisson process: clearly

$$\Psi_{t_n} = \sum_{k \ge 1} \delta_{(t_{n+k} - t_n)e^{\alpha t_n}},$$

and provided that, up to a multiplicative constant, $e^{\alpha t_k}/k$ converges to W, the point process $\sum_{k\geq 1} \delta_{n(t_{n+k}-t_n)}$ should converge to a Poisson random variable with a deterministic parameter. In this case the probability that a ball falls into the *n*th bin which is given by

$$P_n = e^{-\rho t_{n-1}} (1 - e^{-\rho (t_n - t_{n-1})}),$$

has therefore the following asymptotic behavior

$$P_n \sim n^{-\rho/\alpha - 1} W^{\rho/\alpha} F_n,$$

where (F_i) are i.i.d. exponential random variables with mean ρ . In the Bellman-Harris case, following Athreya and Kaplan [**AK76**], it is possible to show that Wand (F_i) are independent, so that in this case, the asymptotic behavior of (P_n) is exactly the same as in the case of a Yule process. One can conjecture that this independence property still holds in the general case.

The main obstacle to generalize the results of this chapter, even in the Bellman-Harris case, is that although W and (F_i) are independent, t_{n-1} and $t_n - t_{n-1}$ are not independent. In the proof of Proposition 1, this independence plays a crucial role, it has therefore to be generalized to variables which are only asymptotically independent. Additionally, since the heavy tail property of the limiting variable $W_{\infty}^{-\rho}$ is also true in the general case, see e.g., Liu [Liu01], a similar rare events phenomenon to the one described in Section 5 is plausible in this case.

General Balls and Yule Process. When the underlying branching process is changed, the above discussion suggests that the asymptotic behavior of the sequence (P_n) remains essentially the same as for a Yule process. The situation changes significantly when the law of the location X of a ball is changed, in this case with the same notations as before for the Yule process,

$$P_n = \mathbb{P}(t_{n-1} < X \le t_{n-1} + E_n/n).$$

The tail distribution of X then plays a key role. Consider for instance a power law, i.e., $\mathbb{P}(X \ge x)$ behaves as $\delta x^{-\beta}$ for some β and $\delta > 0$: then

$$P_{n+1} \sim t_n^{-\beta} - (t_n + E_{n+1}/(n+1))^{-\beta} \sim \frac{\beta \delta E_{n+1}}{n t_n^{\beta+1}} \sim \frac{\beta \delta E_{n+1}}{n (\log n)^{\beta+1}},$$



FIGURE 1. Renewal bins and balls problem on \mathbb{R}_+ : S_n is a renewal process, the balls (E_n) are i.i.d. exponential random variables.

and it can be seen that the random variable W_{∞} may not play a role anymore in the asymptotic behavior of the system.

7. Comparison with the Renewal Case

In a series of papers [Gne04, GINR09, GIR08], Gnedin *et al.* look at the case where the point process used to divide \mathbb{R}_+ into random intervals is a renewal process (S_n) , instead of the split times (t_n) of a Yule process (see Figure 1). If X is the law of a step of S, then the random probability distribution (P_n) describing how balls are thrown is given by

(13)
$$P_n = (1 - W_n) \prod_{i=1}^{n-1} W_i,$$

where (W_n) are i.i.d. random variables on (0, 1) with common distribution $W = e^{-X}$. In contrast to our case where (P_n) essentially decays as a power law, here the deterministic analog is a geometric distribution, which gives raise to different behaviors. We discuss in this section the fact that, as far as the locations of empty bins are concerned, there are different areas of interest in the Yule case, and only one in the renewal case. To illustrate this idea, we begin by proving the simple following lemma for the Yule case, i.e., the case studied so far:

LEMMA 7.1. Let L_n be the index of the last bin which contains at least one ball when n balls are thrown. Then the sequence $(L_n/n^{1/\rho})$ converges in distribution to a random variable L_{∞} , positive and finite almost surely.

PROOF. The proof relies on building L_n in a probability space where, rescaled by $n^{1/\rho}$, it converges almost surely to a finite random variable $L_{\infty} > 0$. Let (t_n) be the split times of a Yule process, (E_n) i.i.d. exponential random variables with mean $1/\rho$, and $S_n = \sum_{1}^{n} E_k/k$. Define Λ_n as the index such that $t_{\Lambda_n} \leq S_n < t_{\Lambda_n+1}$: since S_n has the same law as the maximum of n i.i.d. exponential random variables, Λ_n and L_n have the same law, and so it is enough to show that $\Lambda_n/n^{1/\rho}$ converges almost surely. It has already been seen that $M_n = S_n - (1/\rho) \log n$ and $M'_n = t_n - \log n$ both converge almost surely to M_{∞} and M'_{∞} respectively: by definition of Λ_n ,

$$0 \le S_n - t_{\Lambda_n} = M_n + (1/\rho) \log n - M'_{\Lambda_n} - \log \Lambda_n \le t_{\Lambda_n + 1} - t_{\Lambda_n}$$

Since Λ_n goes to $+\infty$ with probability one and $t_{n+1} - t_n$ goes to 0 with probability one as well, one readily gets from the previous inequalities that almost surely,

$$\lim_{n \to +\infty} \log \left(n^{1/\rho} / \Lambda_n \right) = M'_{\infty} - M_{\infty},$$

and the result immediately follows. For the sake of completeness, note that we know the distribution of L_{∞} , since M_{∞} and M'_{∞} are independent exponential random variables.

Range of Indexes in the Yule Case. Theorem 4.1 shows that the first empty bin has an index of order $n^{1/(\rho+2)}$, and states more precisely that asymptotically, the number of empty bins with indices lying in the range $(xn^{1/(\rho+2)}, yn^{1/(\rho+2)})$ for x < y is finite. The last empty bin has an index of order $n^{1/\rho}$ by the previous lemma, and Proposition 3.2 suggests (it can be rigorously proved) that if $1/(\rho+2) < \alpha$, then the number of empty bins with indices lying in the range $(xn^{\alpha}, yn^{\alpha})$ for x < yscales like $n^{1/(\alpha(\rho+2)-1)}$ for $\alpha \leq 1/(\rho+1)$, and like n^{α} for $\alpha \geq 1/(\rho+1)$. Hence there are infinitely many empty bins between the first empty bin and the last non-empty one, and the density of these bins smoothly increases in between.

Range of Indexes in the Renewal Case. In the renewal case, i.e., when the random probability distribution (P_n) is given by (13), then there is only one range of interest: for the sake of simplicity, and unless otherwise specified, we state the results of Gnedin *et al.* [GINR09] in the simplest case where both $\mu = \mathbb{E}X$ and $\mathbb{E}(X^2)$ are finite. Under this condition, the indices ν_n of the first empty bin and L_n of the last non-empty one are sharply centered around $(1/\mu) \log n$: both

$$\frac{\nu_n - (1/\mu)\log n}{\sqrt{\log n}} \quad \text{and} \quad \frac{L_n - (1/\mu)\log n}{\sqrt{\log n}}$$

converge in distribution to normal random variables; thus the first empty bin and the last empty one are both of order of log n. Concerning the *number* of empty bins with index smaller than L_n , Gnedin *et al.* prove that this number converges, without rescaling, towards a finite random variable. In the case $\mathbb{E}X = +\infty$ of infinite expectation, this number converges to 0 because $\nu_n - L_n$ converges (in distribution) to 1: this means that there is no whole, i.e., all bins are non-empty, and then all bins are empty, which is quite an unexpected property.

These results suggest that not only are L_n and ν_n of the same order, but they should be very close to each other. Although Gnedin *et al.* do not prove this result, it readily follows from the following very nice construction, which they introduce in **[GIR08]** and that we present here. The key idea is that in the renewal case, there exists a bins and balls problem which represents what happens asymptotically around the last non-empty bin. It readily follows from this asymptotic problem that the first empty bin and the last non-empty one are only separated by a finite number of bins.

Asymptotic Problem. So far, balls were thrown on \mathbb{R}_+ which was divided into random intervals. By applying the map $x \mapsto e^{-x}$ to both the intervals and the balls, this problem can be cast on (0, 1) — see Figure 2. With this transformation, the *n* balls become *n* i.i.d. points on (0, 1) uniformly distributed; and the random intervals (S_n, S_{n+1}) of \mathbb{R}_+ become random intervals (Q_{n+1}, Q_n) of (0, 1), where Q_n is defined by

$$Q_0 = 1$$
 and $Q_n = P_{n+1} - P_n = \prod_{k=1}^n W_k = e^{-S_n}, \ n \ge 1.$



FIGURE 2. Renewal bins and balls problem on (0, 1) obtained by applying the map $x \mapsto e^{-x}$ to the renewal bins and balls problem on \mathbb{R}_+ : $Q_n = e^{-S_n}$, the balls (U_n) are i.i.d. uniform random variables.



FIGURE 3. Asymptotic renewal bins and balls problem on \mathbb{R}_+ : $\widetilde{Q}_k = e^{-\widetilde{S}_k}$ with \widetilde{S}_n the stationary renewal process on \mathbb{R} with step distribution X, the balls (σ_n) are the point of a Poisson process with intensity one.

As usual, we consider that a ball falls in the *n*th interval if it falls in the interval (Q_n, Q_{n-1}) .

Of course, the problem on (0, 1) can be considered on any finite interval (0, a), by just scaling everything by a > 0: the intervals are then of the form (aQ_{n+1}, aQ_n) and the balls correspond to a sequence (aU_n) where (U_n) are i.i.d. uniform on (0, 1). Consider now that n balls are thrown on (0, 1), and consider the specific scaling factor a = n: balls then form a sequence $(nU_k, 1 \le k \le n)$, and intervals are defined by the points $(e^{-S_k + \log n}, k \ge 1)$. Quite remarkably, both sequences converge under mild assumptions on X as n goes to infinity.

On the one hand, it is well-known that the sequence $(nU_k, 1 \le k \le n)$ converges in distribution to a Poisson process on \mathbb{R}_+ of intensity one.

On the other hand, the renewal theorem states that as t goes to infinity, the sequence $(S_n - t, n \ge 1)$ converges in distribution to a stationary renewal process on \mathbb{R} ; see for instance Robert [**Rob03**, Proposition 1.28]. In particular, the sequence $(nQ_k, k \ge 0)$ converges in distribution, as n goes to infinity, to a sequence $(e^{-\tilde{S}_k}, -\infty < k < +\infty)$ where \tilde{S} is the stationary renewal process on \mathbb{R} with step distribution X.

Finally, the asymptotic bins and balls problem of Figure 3 arises when the two convergences are considered simultaneously: balls are given by a Poisson process on \mathbb{R}_+ with intensity one, whereas bins are defined by the open intervals between the points of the sequence $(e^{-\tilde{S}_n}, -\infty < n < +\infty)$ with \tilde{S} a stationary renewal process on \mathbb{R} . In this asymptotic picture, the first non-empty interval (starting from the left) corresponds to the last non-empty interval.

Intuitively, any quantity of interest for the renewal bins and balls problem can be read off this limiting problem. For instance, Gnedin *et al.* show that the number of balls in the last non-empty interval for the renewal bins and balls problem converges in distribution to the number of balls in the first non-empty interval in the limiting problem; it can be proved similarly that in the limiting problem, there is a finite index η — where bins are labelled from left to right starting from the first non-empty one — such that there is no empty bin after this index (observe that the intervals are larger and larger, whereas the balls are regularly spaced). Then the difference between the index of the first empty bin and the index of the last non-empty one in the renewal bins and balls problem converges to η .

Appendix A. Proof of Proposition 5.1

This section is devoted to the proof of Proposition 5.1. The notations of Section 5 are used. For $n \ge 1$, one denotes by $k_n = \lfloor x\phi(n) \rfloor$, it is assumed that $\phi(n) \ll n^{\alpha_0}$ for some $\alpha_0 < 1/(\rho + 1)$.

For $0 \le p \le 1$ and $n \ge 1$, the elementary inequality

$$e^{-np} - (1-p)^n \Big| \le \frac{p^2}{2} n e^{-np} \le \frac{2e^2}{n}$$

and the relation $\mathbb{E}\left(\mathcal{N}_{n}^{\phi}([0,x])\right) = \mathbb{E}[(1-P_{1})^{n}] + \dots + \mathbb{E}[(1-P_{\lfloor x\phi(n) \rfloor})^{n}]$ give directly the following lemma.

LEMMA A.1. For a non-decreasing function ϕ , $x \ge 0$, and $n \ge 1$, then

$$\left| \mathbb{E} \left(\mathcal{N}_n^{\phi}([0,x]) \right) - \sum_{i=1}^{k_n} \mathbb{E} \left(e^{-nP_i} \right) \right| \le 2e^2 \frac{k_n}{n}.$$

LEMMA A.2. There exists $\eta > 0$ such that

$$\mathbb{E}\left(\sup_{n\geq 1}W_n^{-\eta}\right) < +\infty.$$

PROOF. For $n \ge 1$, then $W_n = \exp(-V_{n-1})\exp(-(H_{n-1} - \log n))$ where $V_n =$ $t_n - H_n$ and (H_n) is the harmonic sequence. For $0 < \eta < 1/2$, it is easy to check that

 $\sup \left\{ \mathbb{E} \left(\exp(2\eta V_n) \right) : n \ge 0 \right\} < +\infty.$

The sequence (V_n) being a martingale, Doob's inequality that for any $n \ge 0$

$$\mathbb{E}\left(\sup_{n\geq 1}e^{2\eta V_n}\right)\leq 4\sup_{n\geq 1}\mathbb{E}\left(e^{2\eta V_n}\right).$$

The result is proved.

LEMMA A.3. If F is an exponential random variable with mean ρ independent of W_i then

$$\lim_{n \to \infty} \left| \mathbb{E} \left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1} W_i^{\rho} Z_i} \right) - \mathbb{E} \left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1} W_i^{\rho} F} \right) \right| = 0.$$

PROOF. Let y > 0 and $Z = (1 - e^{-yF})/y$. Then, trite calculations with the exponential distribution give that, for any $\beta > 0$,

$$0 \leq \mathbb{E}(e^{-\beta Z}) - \mathbb{E}(e^{-\beta F}) \leq e^{-(1+\rho\beta)/(2y)} + \frac{2y}{(1+\rho\beta)^2}$$

For $k \ge 1$ and $1 \le i \le k$, since $Z_i = i \left(1 - e^{-\rho E_i/i}\right)$ where E_i is an exponential random variable with mean 1, by using this relation one obtains that

$$(14) \quad 0 \leq \mathbb{E}\left(\sum_{i=1}^{k} e^{-ni^{-\rho-1}W_{i}^{\rho}Z_{i}}\right) - \mathbb{E}\left(\sum_{i=1}^{k} e^{-ni^{-\rho-1}W_{i}^{\rho}F}\right)$$
$$\leq \mathbb{E}\left(\sum_{i=1}^{k} \exp\left(-\frac{i}{2\rho}(1+\rho ni^{-\rho-1}W_{i}^{\rho})\right)\right) + \mathbb{E}\left(\sum_{i=1}^{k} \frac{2\rho}{i(1+\rho ni^{-\rho-1}W_{i}^{\rho})^{2}}\right).$$

Since $(i^{-\rho-1}W_i^{\rho})$ is a non-increasing sequence, then

$$\begin{split} \sum_{i=1}^{k_n} \exp\left(-\frac{i}{2\rho}(1+\rho n i^{-\rho-1}W_i^{\rho})\right) &\leq \sum_{i=1}^{k_n} \exp\left(-\frac{i}{2\rho}(1+\rho n k_n^{-\rho-1}W_{k_n}^{\rho})\right) \\ &\leq \frac{\exp\left(-(1+\rho n k_n^{-\rho-1}W_{k_n}^{\rho})/(2\rho)\right)}{1-\exp\left(-(1+\rho n k_n^{-\rho-1}W_{k_n}^{\rho})/(2\rho)\right)} \end{split}$$

and this last term converges almost surely to 0 because $k_n \ll n^{1/(\rho+1)}$. The relation

$$\sum_{i=1}^{k_n} \exp\left(-(1+\rho n i^{-\rho-1} W_i^{\rho})\frac{i}{2\rho}\right) \le \sum_{i=1}^{\infty} \exp\left(-\frac{i}{2\rho}\right) < +\infty,$$

gives that the first term in the right hand side of (14) vanishes by Lebesgue's Theorem.

The second term in the right hand side of (14) converges almost surely to 0

$$\sum_{i=1}^{k_n} \frac{1}{i(1+\rho n i^{-\rho-1} W_i^{\rho})^2} \le \frac{1}{(1+\rho n k_n^{-\rho-1} W_{k_n}^{\rho})^2} \sum_{i=1}^{k_n} \frac{1}{i}$$

since $k_n \ll n^{1/(\rho+1)}$. For $\varepsilon > 0$, again by monotonicity,

$$\sum_{i=1}^{k_n} \frac{1}{i(1+\rho n i^{-\rho-1} W_i^{\rho})^2} = \sum_{i=1}^{k_n} \frac{i^{\varepsilon}}{i^{1+\varepsilon}(1+\rho n i^{-\rho-1} W_i^{\rho})^2} \le C_{\varepsilon} \frac{k_n^{\varepsilon}}{(1+\rho n k_n^{-\rho-1} W_{k_n}^{\rho})^2}$$

where C_{ε} is a finite constant. By using some elementary calculations and Lemma A.2, it is not difficult to show that one can choose $\varepsilon > 0$ so that this last term is uniformly bounded in n by an integrable random variable. Again Lebesgue's Theorem shows that the second term in the right of (14) converges to 0. The lemma is proved.

Combined with the last result, the following lemma completes the proof of Proposition 5.1.

LEMMA A.4. If F is an exponential random variable with mean ρ independent of W_i^ρ then

$$\lim_{n \to +\infty} \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_i^{\rho}F}\right) \middle/ \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_{\infty}^{\rho}F}\right) = 1,$$

holds under the assumption of Proposition 5.1 that

$$\sum_{i=1}^{k_n} \mathbb{E}\left(e^{-n\widetilde{P}_i}\right)$$

converges to infinity as n gets large.

Proof. For any $n \ge 1$:

$$\begin{split} \left| \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_i^{\rho}F}\right) - \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_{\infty}^{\rho}F}\right) \right| \\ &= \left| \mathbb{E}\left(\sum_{i=1}^{k_n} \frac{1}{1+\rho ni^{-\rho-1}W_i^{\rho}} - \sum_{i=1}^{k_n} \frac{1}{1+\rho ni^{-\rho-1}W_{\infty}^{\rho}}\right) \right| \\ &\leq \sum_{i=1}^{k_n} \mathbb{E}\left(\frac{\rho ni^{-\rho-1}|W_{\infty}^{\rho} - W_i^{\rho}|}{(1+\rho ni^{-\rho-1}W_i^{\rho})(1+\rho ni^{-\rho-1}W_{\infty}^{\rho})}\right) \\ &\leq \sum_{i=1}^{k_n} \mathbb{E}\left(\frac{|W_i^{\rho}/W_{\infty}^{\rho} - 1|}{1+\rho ni^{-\rho-1}W_i^{\rho}}\right). \end{split}$$

hence,

$$\begin{split} \left| \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_i^{\rho}F}\right) - \left| \mathbb{E}\left(\sum_{i=1}^{k_n} e^{-ni^{-\rho-1}W_{\infty}^{\rho}F}\right) \right| \\ & \leq \sum_{i=1}^{k_n} \mathbb{E}\left(\frac{1}{1+\rho ni^{-\rho-1}W_i^{\rho}}\right) \mathbb{E}\left|\frac{W_i^{\rho}}{W_{\infty}^{\rho}} - 1\right|, \end{split}$$

by independence of $W_i^{\rho}/W_{\infty}^{\rho}$ and W_i^{ρ} . The sequence $(W_i^{\rho}/W_{\infty}^{\rho})$ being uniformly integrable, it converges to 1 only almost surely and in L_1 , consequently

$$\sum_{i=1}^{k_n} \mathbb{E}\left(\frac{1}{1+\rho n i^{-\rho-1} W_i^{\rho}}\right) \mathbb{E}\left|\frac{W_i^{\rho}}{W_{\infty}^{\rho}} - 1\right| = o\left(\sum_{i=1}^{k_n} \mathbb{E}\left(\frac{1}{1+\rho n i^{-\rho-1} W_i^{\rho}}\right)\right)$$

since the sequence

$$\left(\sum_{i=1}^{k_n} \mathbb{E}\left(e^{-n\tilde{P}_i}\right)\right) = \left(\sum_{i=1}^{k_n} \mathbb{E}\left(\frac{1}{1+\rho n i^{-\rho-1} W_i^{\rho}}\right)\right)$$

converges to infinity. The proof is completed.

References of Chapter IV

- [AK76] K. B. Athreya and N. Kaplan. Convergence of the Age Distribution in the One-Dimensional Supercritical Age-Dependent Branching Process. Annals of Probability, 4(1):38–50, February 1976.
- [AN72] K. B. Athreya and P. E. Ney. *Branching processes*. Springer, 1972.
- [BHJ92] A. D. Barbour, Lars Holst, and Svante Janson. *Poisson approximation*. The Clarendon Press Oxford University Press, New York, 1992.
- [CF76] Endre Csáki and Antónia Földes. On the first empty cell. *Studia Scientarum Mathematicarum Hungarica*, 11:373–382, 1976.
- [Daw93] Donald A. Dawson. Measure-valued Markov processes, volume 1541 of Lecture Notes in Mathematics. Springer, Berlin, 1993.
- [FM85] Philippe Flajolet and G. Nigel Martin. Probabilistic counting algorithms for data base applications. Journal of Computer and System Sciences, 21:182–209, 1985.
- [GHP07] Alexander V. Gnedin, Ben Hansen, and Jim Pitman. Notes on the occupancy problem with infinitely many boxes: General asymptotics and power laws. *Probability Surveys*, 4:146–171, 2007.
- [GINR09] Alexander V. Gnedin, Alex Iksanov, Pavlo Negadajlov, and Uwe Roesler. The Bernoulli Sieve revisited. To Appear in the Annals of Applied Probability, 2009.
- [GIR08] Alexander V. Gnedin, Alex Iksanov, and Uwe Roesler. Small parts in the Bernoulli sieve. DMTCS proceedings, AI:235–242, 2008.
- [Gne04] Alexander V. Gnedin. The Bernoulli Sieve. *Bernoulli*, 10(1):79–96, 2004.
 [H81] Martin Härnqvist. Limit theorems for point processes generated in a
- general branching process. Advances in Applied Probability, 13:650–668, 1981.
- [HJ08] Hsien-Kuei Hwang and Svante Janson. Local limit theorems for finite and infinite urn models. *Annals of Probability*, 36(3):992–1022, 2008.
- [JK77] Norman L. Johnson and Samuel Kotz. Urn models and their application: an approach to modern discrete probability theory. New York: Wiley, 1977.
- [Kar67] Samuel Karlin. Central limit theorems for certain infinite urn schemes. Journal of Mathematics and Mechanics, 17:373–401, 1967.
- [Kin75] J. F. C. Kingman. The first birth problem for an age-dependent branching process. Annals of Probability, 3(5):790–801, October 1975.
- [KS78] Valentin F. Kolchin and Boris A. Sevastyanov. Random allocations. V. H. Winston & Sons, Washington, D.C., 1978.
- [Liu01] Quansheng Liu. Asymptotic properties and absolute continuity of laws stable by a random weighted mean. *Stochastic Processes and their Applications*, 95:83–107, 2001.
- [Ner81] Olle Nerman. On the convergence of supercritical general (C-M-J) branching processes. Z. Wahrscheinlichkeitstheorie verw. Gebiete, 57:365–395, 1981.
- [Nev72] Jacques Neveu. Martingales à temps discret. Masson et Cie, éditeurs, Paris, 1972.
- [Nev77] Jacques Neveu. Processus Ponctuels, volume 598 of École d'été de probabilités de Saint-Flour (P.-L. Hennequin ed.). Springer-Verlag, Berlin, 1977.

- [Rob03] Philippe Robert. Stochastic Networks and Queues. Stochastic Modelling and Applied Probability Series. Springer-Verlag, New York, 2003. xvii+398 pp.
- [Rud87] Walter Rudin. Real and complex analysis, volume third edition. McGraw-Hill Book Co., New York, 1987.
- [SG07] Walid Saddi and Fabrice Guillemin. Measurement based modeling of eDonkey peer-to-peer file sharing system. In *Proceedings of the International Teletraffic Congress 20*, Ottawa, June 2007.
- [SRG08] Florian Simatos, Philippe Robert, and Fabrice Guillemin. A queueing system for modeling a file sharing principle. In *Proceedings of SIGMET-RICS'08*, volume 36, pages 181–192, New York, NY, USA, 2008. ACM.
- [Wil91] David Williams. *Probability with Martingales*. Cambridge University Press, 1991.

Comprehensive list of references

[AFRT06]	Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi. Metastability of CDMA cellular systems. In <i>Mobicom'2006</i> , Los Angeles, September 2006. ACM.
[AFRT08]	Nelson Antunes, Christine Fricker, Philippe Robert, and Danielle Tibi. Stochastic
[AK76]	 K. B. Athreya and N. Kaplan. Convergence of the Age Distribution in the One-Dimensional Supercritical Age-Dependent Branching Process. Annals of Probability, 4(1):38–50. February 1976.
[AK77]	K. B. Athreya and Niels Keiding. Estimation theory for continuous-time branching processes. Sankhya: The Indian Journal of Statistics, 89(A):101–123, 1977.
[AK98]	Krishna B. Athreya and Hye-Jeong Kang. Some limit theorems for positive recur- rent branching Markov chains: I. Advances in Applied Probability, 30(3):711–722, Sentember 1998.
[AKM04]	Guido Appenzeller, Isaac Keslassy, and Nick McKeown. Sizing router buffers. SIG- COMM Computer Communication Review, 34(4):281–292, 2004.
[Ald87]	D. Aldous. Ultimate instability of exponential back-off protocol for acknowledgment- based transmission control of random access communication channels. <i>IEEE Trans-</i> <i>actions on Information Theory</i> , 33(2):219–223, March 1987.
[Als93]	Gerold Alsmeyer. On the Galton-Watson Predator-Prey Process. Annals of Applied Probability, 3(1):198–211, 1993.
[AN72]	K. B. Athreva and P. E. Nev. Branching processes. Springer, 1972.
[AP98]	David Aldous and Jim Pitman. Tree-valued Markov chains derived from Galton-
	Watson processes Ann. Inst. H. Poincaré Probab. Statist., 34(5):637–686, 1998.
[Asm87]	Søren Asmussen. Applied probability and queues. John Wiley & Sons, Chichester, 1987.
[Ath69]	Krishna B. Athreya. On the supercritical one dimensional age dependent branching processes. <i>Annals of Mathematical Statistics</i> , 40(3):743–763, June 1969.
[bAGP ⁺ 05]	Nadia ben Azzouna, Fabrice Guillemin, Stéphanie Poisson, Philippe Robert, Chris- tine Fricker, and Nelson Antunes. Inverting sampled ADSL traffic. In <i>Proceedings of</i> <i>IEEE ICC'05</i> , Seoul, 2005.
[Bar09]	A. D. Barbour. Coupling a branching process to an infinite dimensional epidemic process. To appear in ESAIM, 2009.
[BBP04]	T. Bonald, S. C. Borst, and A. Proutière. How Mobility Impacts the Flow-Level Performance of Wireless Data Systems. In <i>Proceedings of IEEE INFOCOM 2004</i> , 2004.
[BFOBR02]	N. Benameur, S. Ben Fredj, S. Oueslati-Boulahia, and J. W. Roberts. Quality of service and flow level admission control in the internet. <i>Computer Networks</i> , 40(1):57–71, 2002.
[BFP09]	Thomas Bonald, Mathieu Feuillet, and Alexandre Proutière. Is the "Law of the Jun- gle" sustainable for the Internet? In <i>Proceedings of IEEE INFOCOM 2009</i> , April 2009.
[BGM ⁺ 05]	Lawrence Brown, Noah Gans, Avishai Mandelbaum, Anat Sakov, Haipeng Shen, Sergey Zeltyn, and Linda Zhao. Statistical Analysis fo a Telephone Call Center: A Queueing-Science Perspective. <i>Journal of the American Statistical Association</i> , 100(469):36–50, March 2005.
[BH52]	Richard Bellman and Theodore Harris. On Age-Dependent Binary Branching Processes. <i>The Annals of Mathematics</i> , 55(2):280–295, March 1952.

[BHJ92]	A. D. Barbour, Lars Holst, and Svante Janson. <i>Poisson approximation</i> . The Claren-
[BHP06]	A. R. Bharambe, C. Herley, and V. N. Padmanabhan. Analyzing and Improving a
()	BitTorrent Networks Performance Mechanisms. In Proceedings of IEEE INFOCOM
	2006, pages 2884–2895, Barcelona, Spain, April 2006.
[Bit]	Bittorrent.org, http://bittorrent.org.
$[BMM^+08]$	Thomas Bonald, Laurent Massoulié, Fabien Mathieu, Diego Perino, and Andrew
	SICMETRICS'08 pages 325 336 New York NV USA 2008 ACM
[BMP08]	Charles Bordenave, David McDonald, and Alexandre Proutière. Performance of ran-
[2 00]	dom medium access control, an asymptotic approach. In <i>Proceedings of SIGMET-</i> <i>RICS'08</i> , pages 1–12. New York, NY, USA, 2008. ACM.
[Bor05]	Sem Borst. User level performance of channel-aware scheduling algorithms in wireless
	data networks. $I\!EEE/ACM$ Transactions on Networking, 13(3):636–647, June 2005.
[BP02]	Thomas Bonald and Alexandre Proutière. Insensitivity in processor-sharing networks.
	Performance Evaluation, 49(1-4):193-209, September 2002.
[BP03]	I nomas Bonaid and Alexandre Proutiere. Insensitive Bandwidth Sharing in Data Notworks, Oucucing Systems: Theory and Applications 44(1):60, 100, May 2003
[BPH06]	Sem Borst, Alexandre Proutière, and Nidhi Hedge. Capacity of wireless data networks
[]	with intra- and inter-cell mobility. In Proceedings of IEEE INFOCOM 2006, pages
	1058–1069, Barcelona, Spain, April 2006.
[Bra94a]	Maury Bramson. Instability of FIFO Queueing Networks. Annals of Applied Proba-
	<i>bility</i> , 4(2):414–431, 1994.
[Bra94b]	Maury Bramson. Instability of FIFO Queueing Networks with Quick Service Times.
[Bra96a]	Maury Bramson, Convergence to equilibria for fluid models of FIFO queueing net-
[Diaboa]	works. Queueing Systems: Theory and Applications, 22(1–2):5–45, 1996.
[Bra96b]	Maury Bramson. Convergence to equilibria for fluid models of head-of-the-line pro-
	portional processor sharing queueing networks. Queueing Systems: Theory and Ap-
	plications, 23(1–4):1–26, 1996.
[Bra08]	Maury Bramson. Stability of queueing networks, volume 1950 of Lecture Notes in Mathematica, Springer, Barlin, 2008, Lectures from the 26th Drobability Summer
	School held in Saint-Flour. July 2–15, 2006.
[Bur81]	David Y. Burman. Insensitivity in Queueing Systems. Advances in Applied Probabil-
. ,	<i>ity</i> , 13(4):846–859, 1981.
[CF76]	Endre Csáki and Antónia Földes. On the first empty cell. Studia Scientarum Math-
	ematicarum Hungarica, 11:373–382, 1976.
[CMP09]	D. Chafai, F. Malrieu, and K. Paroux. On the long time behavior of the ICP window size process. According Stochastic Processes and their Applications, 2000
[CN04]	F. Clévenot and P. Nain. A simple fluid model for the analysis of the squirrel peer-
[]	to-peer caching system. In Proceedings of IEEE INFOCOM 2004, volume 1, pages
	86–95, Hong Kong, March 2004.
[CP07]	Pietro Caputo and Gustavo Posta. Entropy dissipation estimates in a zero-range
	dynamics. Probability Theory and Related Fields, 139:65–87, September 2007.
	Hong Chen and David D. Yao. Fundamentals of queueing networks. Springer-Verlag,
	and Applied Probability.
[Dai95]	J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified
. ,	approach via fluid limit models. Annals of Applied Probability, 5(1):49–77, February
	1995.
[Daw93]	Donald A. Dawson. Measure-valued Markov processes, volume 1541 of Lecture Notes
[DFT0e]	in Mathematics. Springer, Berlin, 1993.
[DF 100]	ing of wireless relay networks. <i>IEEE Transactions on Information Theory (joint issue</i>
	with IEEE/ACM Transactions on Networking), 52(6):2756–2761, 2006.
[DGR02]	Vincent Dumas, Fabrice Guillemin, and Philippe Robert. A Markovian analysis of
	Additive-Increase Multiplicative-Decrease (AIMD) algorithms. Advances in Applied
	Probability, 34(1):85–111, 2002.

[DHR00]	Jean-François Dantzer, Mostapha Haddani, and Philippe Robert. On the stability of a bandwidth packing algorithm. <i>Probability in the Engineering and Informational</i> <i>Sciences</i> , 14(1):57–79, 2000.
[DLS01]	Bogdan Doytchinov, John Lehoczky, and Steven Shreve. Real-time queues in heavy traffic with earliest-deadline-first queue discipline. <i>Annals of Applied Probability</i> , 11(2):332–378, 2001.
[DPM07]	Trang Dinh Dang, R. Pereczes, and S. Molnar. Modeling the population of file-sharing peer-to-peer networks with branching processes. <i>Computers and Communications</i> , 2007. ISCC 2007. 12th IEEE Symposium on, pages 809–815, 2007.
[Dum97]	Vincent Dumas. A Multiclass Network with Non-Linear, Non-Convex, Non-Monotonic Stability Conditions. <i>Queueing Systems: Theory and Applications</i> , 25(1–4):1–43, 1997.
[DW94]	P. Dupuis and R. J. Williams. Lyapounov functions for semimartingale reflecting Brownian motions. <i>Annals of Applied Probability</i> , 22(2):680–702, 1994.
[EK86]	Stewart N. Ethier and Thomas G. Kurtz. <i>Markov Processes, Characterization and Convergence</i> . Wiley Series in Probability and Mathematical Statistics, 1986.
[Erl09]	A. K. Erlang. Sandsynlighedsregning og Telefonsamtaler. Nyt Tidsskrift for Matematik B, 20:33, 1909.
[Erl17]	A. K. Erlang. Løsning af nogle Problemer fra Sandsynlighedsregningen af Betydning for de automatiske Telefoncentraler. <i>Elektroteknikeren</i> , 13:5, 1917.
[FDTT07]	M. Franceschetti, O. Dousse, D. N. C. Tse, and P. Thiran. Closing the gap in the capacity of random wireless networks via percolation theory. <i>IEEE Transactions on Information Theory</i> , 53(3):1009–10018, 2007.
[FIMM91]	G. Fayolle, I.A. Ignatyuk, V.A. Malyshev, and M.V. Men'shikov. Random walks in two-dimensional complexes. <i>Queueing Systems, Theory and Applications</i> , 9(3):269–300, 1991.
[FJ93]	Sally Floyd and Van Jacobson. Random early detection gateways for congestion avoidance. <i>IEEE/ACM Transactions on Networking</i> , 1(4):397–413, 1993.
[FM85]	Philippe Flajolet and G. Nigel Martin. Probabilistic counting algorithms for data base applications. <i>Journal of Computer and System Sciences</i> , 21:182–209, 1985.
[FRT99]	Christine Fricker, Philippe Robert, and Danielle Tibi. On the rates of convergence of Erlang's model. <i>Journal of Applied Probability</i> , 36(4):1167–1184, 1999.
[GFS ⁺ 03]	Z. Ge, D. R. Figueiredo, Jaiswal Sharad, J. Kurose, and D. Towsley. Modeling peer- peer file sharing systems. In <i>Proceedings of IEEE INFOCOM 2003</i> , volume 3, pages 2188–2198, San Francisco, CA, USA, March 2003.
[GHP07]	Alexander V. Gnedin, Ben Hansen, and Jim Pitman. Notes on the occupancy problem with infinitely many boxes: General asymptotics and power laws. <i>Probability Surveys</i> , 4:146–171, 2007.
[GINR09]	Alexander V. Gnedin, Alex Iksanov, Pavlo Negadajlov, and Uwe Roesler. The Bernoulli Sieve revisited. To Appear in the <i>Annals of Applied Probability</i> , 2009.
[GIR08]	Alexander V. Gnedin, Alex Iksanov, and Uwe Roesler. Small parts in the Bernoulli sieve. <i>DMTCS proceedings</i> , AI:235–242, 2008.
[GK00]	P. Gupta and P. R. Kumar. The capacity of wireless networks. <i>IEEE Transactions</i> on <i>Information Theory</i> , 46(2):388–404, March 2000.
[GMR06]	Christos Gkantsidis, John Miller, and Pablo Rodriguez. Anatomy of a P2P Content Distribution system with Network Coding. In <i>Proceedings of the 5th International Workshop on Peer-to-Peer Systems</i> , Santa Barbara, CA, USA, February 27–28 2006.
[Gne04]	Alexander V. Gnedin. The Bernoulli Sieve. <i>Bernoulli</i> , 10(1):79–96, 2004.
[GPW02]	H. Christian Gromoll, Amber L. Puha, and Ruth J. Williams. The fluid limit of a heavily loaded processor sharing queue. <i>Annals of Applied Probability</i> , 12(3):797–859, 2002.
[Gri99]	Geoffrey Grimmett. Percolation. Springer, 1999.
[GT01]	M. Grossglauser and D. Tse. Mobility increases the capacity of ad-hoc wireless networks. In <i>Proceedings of IEEE INFOCOM 2001</i> , volume 3, pages 1360–1369, 2001.
[GV04]	A. T. Gai and L. Viennot. Broose: A practical distributed hashtable based on the de-bruijn topology. In <i>International Conference on Peer-to-Peer Computing (P2P)</i> , Zurich, August 2004.

[HÏ81]	Martin Härnqvist. Limit theorems for point processes generated in a general branch- ing process. Advances in Applied Probability, 13:650–668, 1981.
[HJ08]	Hsien-Kuei Hwang and Svante Janson. Local limit theorems for finite and infinite urn models. Annals of Probability, 36(3):992–1022, 2008.
[JK77]	Norman L. Johnson and Samuel Kotz. Urn models and their application: an approach to modern discrete probability theory. New York: Wiley, 1977.
[Kar67]	Samuel Karlin. Central limit theorems for certain infinite urn schemes. Journal of Mathematics and Mechanics, 17:373–401, 1967.
[Kel86]	F. P. Kelly. Blocking probabilities in large circuit-switched networks. Advances in Applied Probability, 18(2):473–505, June 1986.
[Kel91]	F. P. Kelly. Loss networks. The Annals of Applied Probability, 1(3):319–378, 1991.
[Kel94]	F. P. Kelly. <i>Reversibility and Stochastic Networks</i> . Wiley, Chichester, 1979, reprinted 1987, 1994.
[KH95]	Raymond Knopp and Pierre A. Humblet. Information capacity and power control in single-cell multiuser communications. In <i>ICC'95</i> , <i>IEEE International Conference on Communications</i> . June 18–22, 1995 – Seattle, USA, June 1995.
[Kin75]	J. F. C. Kingman. The first birth problem for an age-dependent branching process. Annals of Probability, 3(5):790-801, October 1975.
[Kin78]	J. F. C. Kingman. Random partitions in population genetics. Proceedings of the Royal Society. London. Series A. Mathematical, Physical and Engineering Sciences, 361(1704):1-20, 1978.
[KKS09]	G. Kesidis, T. Konstantopoulos, and P. Sousi. A stochastic epidemiological model and a deterministic limit for BitTorrent-like peer-to-peer file-sharing networks. <i>Springer</i> <i>Lecture Notes in Computer Science</i> , 5425:26–36, 2009
[KS78]	Valentin F. Kolchin and Boris A. Sevastyanov. <i>Random allocations</i> . V. H. Winston & Sons, Washington, D.C. 1978.
[Lam67]	John Lamperti. The Limit of a Sequence of Branching Processes. Probability Theory and Related Fields. 7(4):271–288. August 1967.
[Les09]	Lasse Leskelä. Stochastic Relations of Random Variables and Processes. <i>Journal of Theoretical Probability</i> , to appear, 2009.
[Liu01]	Quansheng Liu. Asymptotic properties and absolute continuity of laws stable by a random weighted mean. <i>Stochastic Processes and their Applications</i> , 95:83–107, 2001.
[Loy62]	R. M. Loynes. The stability of a queue with non-independent inter-arrival and service times. <i>Mathematical Proceedings of the Cambridge Philosophical Society</i> , 58(03):497–520, 1962.
[Mal93]	V. A. Malyshev. Networks and dynamical systems. <i>Advances in Applied Probability</i> , 25(1):140–175, 1993.
[MR99]	L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. In <i>Proceedings of IEEE INFOCOM 1999</i> , volume 3, pages 1395–1403 vol.3, March 1999.
[MT08]	Laurent Massoulié and Andrew Twigg. Rate-optimal schemes for Peer-to-Peer live streaming. <i>Performance Evaluations</i> , 65(11-12):804–822, 2008.
[MV05]	Laurent Massoulié and Milan Vojnović. Coupon replication systems. In <i>Proceedings</i> of SIGMETRICS'05, volume 33, pages 2–13, Banff, Alberta, Canada, June 2005.
[Ner81]	Olle Nerman. On the convergence of supercritical general (C-M-J) branching processes. Z. Wahrscheinlichkeitstheorie verw. Gebiete, 57:365–395, 1981.
[Nev72]	Jacques Neveu. Martingales à temps discret. Masson et Cie, éditeurs, Paris, 1972.
[Nev77]	Jacques Neveu. Processus Ponctuels, volume 598 of École d'été de probabilités de Saint-Flour (PL. Hennequin ed.). Springer-Verlag, Berlin, 1977.
[Nev86]	J. Neveu. Erasing a branching tree. Advances in Applied Probability, (suppl.):101–108, 1986.
[NQP08]	R. Núñez-Queija and B. J. Prabhu. Scaling laws for file dissemination in P2P networks with random contacts. In <i>Proceedings of IWQoS</i> , 2008.
[PF95]	Vern Paxson and Sally Floyd. Wide-area traffic: The failure of poisson modeling. <i>IEEE/ACM Transactions on Networking</i> , 3:226–244, 1995.
[PGES05]	Johan Pouwelse, Pawel Garbacki, Dick Epema, and Henk Sips. The Bittorrent P2P File-sharing System: Measurements and Analysis. In <i>Proceedings of IPTPS'05</i> , Ithaca, NY, February 2005.

Comprehensive list of references

166

[PWMC08]	Nadim Parvez, Carey Williamson, Anirban Mahanti, and Niklas Carlsson. Analysis of bittorrent-like protocols for on-demand stored media streaming. In <i>Proceedings of SIGMETRICS'08</i> , pages 301–312, New York, NY, USA, 2008. ACM.
[QS04]	Dongyu Qiu and R. Srikant. Modeling and performance analysis of bittorrent-like peer-to-peer networks. In <i>Proceedings of SIGCOMM'04</i> , pages 367–378, New York, NY, USA, August 2004. ACM.
[Rob03]	Philippe Robert. <i>Stochastic Networks and Queues</i> . Stochastic Modelling and Applied Probability Series. Springer-Verlag, New York, 2003. xvii+398 pp.
[RS92]	A. N. Rybko and A. L. Stolyar. On the ergodicity of random processes that describe the functioning of open queueing networks. <i>Problems on Information Transmission</i> , 28(3):3–26, 1992.
[RS09]	Philippe Robert and Florian Simatos. Occupancy schemes associated to Yule pro- cesses. Advances in Applied Probability, 41(2):600–622, June 2009.
[Rud87]	Walter Rudin. Real and complex analysis, volume third edition. McGraw-Hill Book Co., New York, 1987.
[RW87]	L. C. G. Rogers and David Williams. Diffusions, Markov processes, and martingales. Vol. 2: Itô calculus. John Wiley & Sons Inc., New York, 1987.
[SAV06]	Riikka Susitaival, Samuli Aalto, and Jorma Virtamo. Analyzing the dynamics and resource usage of P2P file sharing by a spatio-temporal model. <i>International Workshop on P2P for High Performance Computational Sciences (P2P-HPCS'06) in conjunction with ICCS</i> , 2006.
[Sch77]	R. Schassberger. Insensitivity of Steady-State Distributions of Generalized Semi- Markov Processes. Part I. <i>The Annals of Probability</i> , 5(1):87–99, 1977.
[Sen70]	E. Seneta. On the supercritical branching process with immigration. <i>Mathematical Biosciences</i> , 7:9–14, 1970.
[SG07]	Walid Saddi and Fabrice Guillemin. Measurement based modeling of eDonkey peer- to-peer file sharing system. In <i>Proceedings of the International Teletraffic Congress</i> 20, Ottawa, June 2007.
[SHM07]	S. Sanghavi, B. Hajek, and L. Massoulié. Gossiping with Multiple Messages. In <i>Proceedings of IEEE INFOCOM 2007</i> , pages 2135–2143, May 2007.
[SRG08]	Florian Simatos, Philippe Robert, and Fabrice Guillemin. A queueing system for modeling a file sharing principle. In <i>Proceedings of SIGMETRICS'08</i> , volume 36, pages 181–192, New York, NY, USA, 2008. ACM.
[ST09]	Florian Simatos and Danielle Tibi. Spatial homogenization in a stochastic network with mobility. <i>Annals of Applied Probability</i> , to appear, 2009.
[Tak69]	Lajos Takács. On Erlang's formula. Annals of Mathematical Statistics, 40(1):71–78, 1969.
[TDG08]	Dominique Tschopp, Suhas Diggavi, and Matthias Grossglauser. Hierarchical rout- ing over dynamic wireless networks. <i>SIGMETRICS Performance Evaluation Review</i> , 36(1):73–84, 2008.
[TWN06]	Ye Tian, Di Wu, and Kam Wing Ng. Modeling, Analysis and Improvement for BitTorrent-Like File Sharing Networks. In <i>Proceedings of IEEE INFOCOM 2006</i> , pages 2921–2931, Barcelona, Spain, April 2006.
[Wal09]	N. S. Walton. Proportional Fairness and its Relationship with Multi-class Queueing Networks. Preprint available on arXiv, 2009.
[Wil91]	David Williams. Probability with Martingales. Cambridge University Press, 1991.
[Yas83]	S. F. Yashkov. A Derivation of Response Time Distribution for an M/G/1 Processor- Sharing Queue. <i>Problems of Control and Information Theory</i> , 12(2):133–148, 1983.
[YdV04]	Xiangying Yang and Gustavo de Veciana. Service capacity of peer to peer networks. In <i>INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer</i> and Communications Societies, volume 4, pages 2242–2252, 2004
[YdV06]	Xiangying Yang and Gustavo de Veciana. Performance of peer-to-peer networks: service capacity and role of resource sharing policies. <i>Performance Evaluations</i> , 63(3):175–194, 2006.
[ZJM00]	Hui Zang, Jason P. Jue, and Biswanath Mukherjee. A review of routing and wave- length assignment approaches for wavelength-routed optical WDM networks. <i>Optical</i> <i>Networks Magazine</i> , 1(1):47–60, January 2000.