

A Queueing System for Modeling a File Sharing Principle

Florian Simatos, Philippe Robert, Fabrice Guillemin

► **To cite this version:**

Florian Simatos, Philippe Robert, Fabrice Guillemin. A Queueing System for Modeling a File Sharing Principle. ACM Sigmetrics: International Conference on Measurement and Modeling of Computer Systems, Jun 2008, Annapolis, United States. 2008. <inria-00262405v2>

HAL Id: inria-00262405

<https://hal.inria.fr/inria-00262405v2>

Submitted on 11 Mar 2008

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

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

A Queueing System for Modeling a File Sharing Principle

Florian Simatos and Philippe Robert
INRIA, RAP project
Domaine de Voluceau, Rocquencourt
78153 Le Chesnay, France
{Florian.Simatos,Philippe.Robert}@inria.fr

Fabrice Guillemin
Orange Labs
2, Avenue Pierre Marzin
22300 Lannion
Fabrice.Guillemin@orange-ftgroup.com

ABSTRACT

We investigate in this paper the performance of a simple file sharing principle. For this purpose, we consider a system composed of N peers becoming active at exponential random times; the system is initiated with only one server offering the desired file and the other peers after becoming active try to download it. Once the file has been downloaded by a peer, this one immediately becomes a server. To investigate the transient behavior of this file sharing system, we study the instant when the system shifts from a congested state where all servers available are saturated by incoming demands to a state where a growing number of servers are idle. In spite of its apparent simplicity, this queueing model (with a random number of servers) turns out to be quite difficult to analyze. A formulation in terms of an urn and ball model is proposed and corresponding scaling results are derived. These asymptotic results are then compared against simulations.

Categories and Subject Descriptors

C.4 [Computer Systems Organization]: Performance of Systems—*modeling techniques, performance attributes*

General Terms

Queueing Systems, Transient Analysis of Markov Processes, File Sharing, Peer to Peer

1. INTRODUCTION

This paper analyzes the performance of a simple file sharing principle during a flash crowd scenario when a popular content becomes available on a peer-to-peer network. It is supposed that a given 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 paper 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 as many servers as peers since each of them can complete 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 this paper.

The analysis performed in this paper 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 [17] 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 paper is precisely to obtain more detailed asymptotics of this transient regime. Except the paper by Yang and de Veciana [17], 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.* [7]. A fluid flow analysis with an underlying Markovian structure is proposed in Clévenot and Nain [5] in order to model the Squirrel peer-to-peer caching system. In Qiu and Srikant [14], the authors directly use a fluid approximation to study the steady state of a peer to peer network, subsequently complemented by diffusion variations around the steady state solutions. In Massoulié and Vojnović [13], the authors study the performance of a file sharing system via a stochastic coupon replication formula-

tion, 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 [14].

The rest of this paper is organized as follows: In Section 2, we describe the system under consideration and some heuristics to study the system are presented. It turns out that the dynamics of the system can be decomposed in two regimes. In the first one, there are almost no empty servers and we establish an analogy with a random urn and ball problem on the real line. By approximating the probability of selecting an urn by its mean value, we analyze in Section 3 the corresponding deterministic urn and ball problem. The analysis for the random urn and ball problem is much more complicated to analyze. The complete analysis is done in [12] and only the main results are summarized in Section 5. In Section 6, we support via simulation the different approximations and heuristics made in this paper to analyze the file sharing system. Concluding remarks are presented in Section 7.

2. MODEL DESCRIPTION

2.1 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.

Exponential distributions. The hypothesis on the distribution on the duration of the time for a peer to be 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 N peers 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 be-

comes 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 not discussed here. As it will be seen, 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 [7, 13, 14], 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 the 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.




Figure 1: Fraction of idle servers: $N=10^6$ and $\rho=5/6$.

2.2 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 as many servers as customers.

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 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 [1] 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 urns 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.

2.3 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 $n \geq 1$, there are by definition exactly n servers. So if we assume, as argued above, that empty servers are negligible during the first regime, $S_n - S_{n-1}$ is well approximated 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 1. *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*

$$T_n = \sum_{k=1}^n \frac{E_k^1}{k}, \quad (1)$$

and E_n^1 being i.i.d. exponential random variables with unit mean.

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 1 and the subsequent discussion below provide strong arguments to support its accuracy. In the definition of the above approximation, it is essential to determine the duration of the first regime, in particular to know whether $S_n \approx T_n$ holds or not.

For instance, one could consider as definition for the duration of the first regime the last time when there are no empty servers. This time is unfortunately not a stopping time and turns out to be much more difficult to study. In Section 6, we shall consider different heuristics for evaluating the length of the first regime. We start the analysis by introducing the index ν defined as follows.

DEFINITION 1. *The duration of the first regime is defined as S_ν , where ν is the first index $n \geq 1$ 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 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 $n < \nu$. \square

We are now able to justify Approximation 1. Indeed, for $n < \nu$, the number of non idle servers is between $n - 2$ and n . For n large, approximately n servers are busy, thus $S_n - S_{n-1}$ is close in distribution to an exponentially distributed random variable with parameter n . During the first time intervals, the number of empty servers is negligible. Indeed,

consider any finite index n , then it is easy to see that the mean number of peers that arrive in the n -th interval is proportional to N . So after the creation of the n -th server, the mean time before the next arrival behaves as $1/N$, and so is very small when N is large. This intuitively shows that the fraction of idle servers is initially negligible, which justifies Approximation 1.

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. URN AND BALL PROBLEM

Denote by $(E_i^\rho, 1 \leq i \leq N)$ an i.i.d. sequence of exponentially distributed random variables 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 urn and ball model on the real line: The interval (T_{n-1}, T_n) is the n -th urn and the variables $(E_i^\rho, 1 \leq i \leq N)$ are the locations of N balls thrown on the real line. The set $\{T_{n-1} \leq E_i^\rho \leq T_n\}$ is simply the event that the i -th ball falls into the n -th urn. Conditionally on the sizes of the urns, i.e., on $\mathcal{T} = (T_n)$, we have that the probability of such an event (which does not depend on i) is

$$P_n = \mathbb{P}(T_{n-1} < E_i^\rho < T_n \mid \mathcal{T}) = e^{-\rho T_{n-1}} \left(1 - e^{-\rho E_n^1/n}\right), \quad (2)$$

where the random variables $E_n^1, n \geq 1$, are independent and exponentially distributed with mean unity.

With the above formulation, we have then to deal with the following urn and ball model:

1. A random probability distribution $\mathcal{P} = (P_n)$ is given (urns with random sizes).
2. N balls are thrown independently according to the probability distribution \mathcal{P} .

It is worth noting that the above urn and ball model has an infinite number of urns. In addition, although urn and ball problems have been widely studied in the literature, our model presents a remarkable feature: For $i \geq 1$, a ball falls into urn i with probability P_i which is a random variable, but conditionally on the sequence (P_n) , this is a classical urn and ball problem. Mathematical results for urn models with random distributions are quite rare. See Kingman [11] and Gnedin *et al.* [8] 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 2. *Let $(E_i^1), i \geq 1$, be independent exponential random variables with parameter 1. Then, for $n \in \mathbb{N}$*

$$T_n = \sum_{k=1}^n \frac{E_k^1}{k} \stackrel{\text{dist.}}{=} \max_{1 \leq k \leq n} E_k^1, \quad (3)$$

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 n -th urn can be written as

$$P_n = \frac{\rho}{n^{\rho+1}} X_{n-1} Z_n, \quad (4)$$

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 $q \geq 1$.

The limiting variable Z_∞ has an exponential distribution with parameter 1 and X_∞ has a Weibull distribution with parameter $1/\rho$,

$$\mathbb{P}(X_\infty \geq x) = e^{-x^{1/\rho}}, \quad x \geq 0. \quad (5)$$

PROOF. Let $E_{(1)} \leq E_{(2)} \leq \dots \leq E_{(n)}$ 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 $E_{(i+1)} - E_{(i)}, i = 0, \dots, n-1$ 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 $Z_n \stackrel{\text{dist.}}{=} n/\rho(1 - \exp(-\rho E_1^1/n))$, 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, $H_n = 1 + 1/2 + \dots + 1/n$. The sequence (M_n) is clearly a martingale, it is bounded in L_2 since

$$\mathbb{E}M_n^2 = \sum_{k=1}^n \frac{\mathbb{E}(E_k^1 - 1)^2}{k^2} = \sum_{k=1}^n \frac{1}{k^2} < +\infty.$$

It therefore converges almost surely. See Williams [16] 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 $x \geq 0$,

$$\mathbb{P}(T_n - \log n \leq x) = (1 - e^{-x - \log n})^n \sim e^{-e^{-x}},$$

as n goes to infinity.

Since

$$X_n = e^{-\rho M_n} e^{\rho(\log(n+1) - H_n)},$$

one gets the almost sure convergence of (X_n) . It is easy to

check that, for $q \geq 0$,

$$\begin{aligned}\mathbb{E}(X_n^q) &= (n+1)^{q\rho} \prod_{i=1}^n \frac{1}{1+q\rho/i} \\ &= (n+1)^{q\rho} \frac{\Gamma(n)}{\Gamma(n+q\rho)} \Gamma(q\rho) \sim \Gamma(q\rho),\end{aligned}\quad (6)$$

when $n \rightarrow \infty$, where Γ is the usual Gamma function, and where the last equivalence easily comes from Stirling's Formula. In particular, for any $q \geq 0$, 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_∞ . \square

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 urn satisfies the following relation

$$\mathbb{E}(P_n) \sim \frac{\rho\Gamma(\rho)}{n^{\rho+1}}. \quad (7)$$

This equivalence suggests the introduction of a deterministic version of the urn and ball problem considered.

4. DETERMINISTIC PROBLEM

4.1 Description

Denote by $\mathcal{Q} = (q_n)$ a probability distribution on \mathbb{N} such that

$$\lim_{n \rightarrow +\infty} n^\delta q_n = \alpha, \quad (8)$$

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 urn. 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 urn and ball 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 urn 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 urn when N balls have been thrown in the random (resp. deterministic) urn and ball problem, and define*

$$\begin{aligned}\nu^R(N) &= \inf\{i \geq 1 : \eta_i^R(N) = 0\}, \\ \nu^D(N) &= \inf\{i \geq 1 : \eta_i^D(N) = 0\}.\end{aligned}\quad (9)$$

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 $\nu'(N) = \inf\{i \geq 1 : \eta_i(N) \leq 1\}$ to be rigorous. In fact, the mathematical analysis of the index of the first empty urn can easily be extended to the first urn that receives less than k balls. 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 4.2, 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 urn 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 1 below shows that the location of the first empty urn is in fact slightly smaller than $N^{1/(\rho+1)}$, i.e., of the order of $(N/\ln N)^{1/(\rho+1)}$. 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 urn 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

$$\mathbb{P}\left(\frac{E_i^1}{i} \leq \frac{i^{\rho+1}}{N} E_0^1\right) = \frac{1}{1 + N/i^{\rho+2}}.$$

When $N \rightarrow \infty$, 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 2 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.

4.2 Asymptotic Analysis

Csáki and Földes [6] 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 urns (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 urns whose index is less than k when N balls have been

thrown. This random variable is formally defined as

$$W_N^k = \sum_{i=1}^k I_{N,i}, \text{ with } I_{N,i} = \mathbb{1}_{\{\eta_i^D(N)=0\}}. \quad (10)$$

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 3. *Assume that the sequence (q_i) is non-increasing. For $x > 0$, if*

$$\kappa_x(N) = \left\lfloor \left(\alpha \delta \frac{N}{\log N} \right)^{1/\delta} \left[1 + \frac{1+\delta}{\delta} \frac{\log \log N}{\log N} + \frac{\log x}{\log N} \right] \right\rfloor, \quad (11)$$

where $\lfloor y \rfloor$ is the integral part of $y > 0$, then

$$\lim_{N \rightarrow +\infty} \mathbb{E} \left(W_N^{\kappa_x(N)} \right) = (\alpha \delta)^{1/\delta} x. \quad (12)$$

PROOF. For $k, N \in \mathbb{N}$

$$\mathbb{E} \left(W_N^k \right) = \sum_{i=1}^k (1 - q_i)^N. \quad (13)$$

For $0 \leq x \leq 1$,

$$0 \leq e^{-Nx} - (1-x)^N \leq x_N (1-x_N)^{N-1},$$

where x_N is the unique solution to the equation $\exp(-Nx) = (1-x)^{N-1}$, since the function $x \rightarrow e^{-Nx} - (1-x)^N$ has a maximum at point x_N . It is easily seen that $Nx_N \leq 2$ (in fact $Nx_N \rightarrow 2$ as $N \rightarrow +\infty$), so that for $N \geq 1$

$$\sup_{0 \leq x \leq 1} \left| e^{-Nx} - (1-x)^N \right| \leq \frac{2}{N}. \quad (14)$$

With this relation, we obtain

$$\left| \mathbb{E} \left(W_N^k \right) - \sum_{i=1}^k e^{-Nq_i} \right| \leq \frac{2k}{N},$$

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 $i \geq 1$. 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 $x \geq 0$.

$$\int_0^k e^{-Nq(u)} du \leq \sum_{i=1}^k e^{-Nq_i} \leq \int_1^{k+1} e^{-Nq(u)} du.$$

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

$$\begin{aligned} & \int_1^{k(N)} e^{-\alpha N u^{-\delta}} du \\ &= \frac{(\alpha N)^{1/\delta}}{\delta} \int_{\alpha N k(N)^{-\delta}}^{\alpha N} e^{-u} u^{-(\delta+1)/\delta} du. \end{aligned} \quad (15)$$

Now let $H(N) = \alpha N k(N)^{-\delta}$ and consider

$$\begin{aligned} & e^{H(N)} H(N)^{(1+\delta)/\delta} \int_{H(N)}^{\alpha N} e^{-u} u^{-(\delta+1)/\delta} du \\ &= \int_{H(N)}^{\alpha N} e^{-(u-H(N))} \left(\frac{H(N)}{u} \right)^{-(\delta+1)/\delta} du \\ &= \int_1^{\alpha N/H(N)} H(N) e^{-H(N)(u-1)} \frac{1}{u^{(\delta+1)/\delta}} du \\ &\sim \int_0^{+\infty} H(N) e^{-H(N)u} \frac{1}{(1+u)^{(\delta+1)/\delta}} du \sim 1, \end{aligned}$$

since $N/H(N) \rightarrow +\infty$ and $H(N) \rightarrow +\infty$ as $N \rightarrow +\infty$. Therefore, an equivalent expression of the integral in the right hand side of Equation (15) has been obtained. Gathering these results, we obtain

$$\begin{aligned} \mathbb{E} \left(W_N^{k(N)} \right) &= \frac{(\alpha N)^{1/\delta}}{\delta} e^{-H(N)} H(N)^{-(1+\delta)/\delta} + o(1) \\ &\sim \frac{1}{\alpha \delta} \frac{k(N)^{1+\delta}}{N} \exp \left(-\alpha N k(N)^{-\delta} \right). \end{aligned} \quad (16)$$

Relation (12) is obtained by taking $k(N) = \kappa_x(N)$. \square

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. *Assume that the sequence (q_i) is non-increasing. For $x > 0$, let κ_x be defined by Equation (11), then*

$$\lim_{N \rightarrow +\infty} \text{Var} \left(W_N^{\kappa_x(N)} \right) / \mathbb{E} \left(W_N^{\kappa_x(N)} \right) = 1. \quad (17)$$

PROOF. For $k \geq 1$, by using Equation (13) (which does not depend on α).

$$\mathbb{E}[W_N^k]^2 = \sum_{1 \leq i, j \leq k} (1 - q_i - q_j + q_i q_j)^N,$$

and

$$\mathbb{E}[(W_N^k)^2] = \mathbb{E}[W_N^k] + \sum_{1 \leq i \neq j \leq k} (1 - q_i - q_j)^N,$$

so that, to prove the equivalence of $\text{Var}(W_N^k)$ and $\mathbb{E}(W_N^k)$, it is sufficient to show that the quantities

$$\sum_{1 \leq i, j \leq k} \left[(1 - q_i - q_j + q_i q_j)^N - (1 - q_i - q_j)^N \right]$$

and

$$\sum_{i=1}^k (1 - 2q_i)^N$$

are negligible with respect to $\mathbb{E}(W_N^k)$. Since we consider $k(N) = \kappa_x(N)$, this amounts to show that these quantities are $o(1)$ by Proposition 4. The second term is the expected number of empty urns for the distribution (\tilde{q}_i) such that $\tilde{q}_i \sim 2\alpha/i^\delta$. Estimate (16) shows that

$$\begin{aligned} \sum_{i=1}^{\kappa_x(N)} (1-2q_i)^N &\sim \frac{1}{2\alpha\delta} \frac{\kappa_x(N)^{1+\delta}}{N} \exp\left(-2\alpha N \kappa_x(N)^{-\delta}\right) \\ &= o\left(\mathbb{E}W_N^{\kappa_x(N)}\right). \end{aligned}$$

By using the fact that for $a \geq b \geq 0$, $a^N - b^N \leq N(a - b)a^{N-1}$, the second term satisfies

$$\begin{aligned} &\sum_{1 \leq i, j \leq k} \left[(1 - q_i - q_j + q_i q_j)^N - (1 - q_i - q_j)^N \right] \\ &\leq N \sum_{1 \leq i, j \leq k} q_i q_j (1 - q_i - q_j + q_i q_j)^{N-1} \\ &= \frac{1}{N} \left(\sum_{i=1}^k N q_i (1 - q_i)^{N-1} \right)^2. \quad (18) \end{aligned}$$

By using a similar method as in the proof of Proposition 3, we obtain the equivalence

$$\begin{aligned} \sum_{i=1}^{k(N)} N q_i (1 - q_i)^{N-1} &\sim \int_1^{k(N)} N q(u) e^{-Nq(u)} du \\ &\sim (\alpha\delta)^{1/\delta} \alpha x \frac{N}{\kappa_x(N)^\delta} = (\alpha\delta)^{1/\delta} x \log N. \end{aligned}$$

This equivalence together with Equation (18) complete the proof of the proposition. \square

THEOREM 1. *Let (q_n) be a non-increasing sequence satisfying Relation (8). For $x > 0$ and $N \in \mathbb{N}$, set*

$$\kappa_x(N) = \left\lceil \left(\alpha\delta \frac{N}{\log N} \right)^{1/\delta} \left(1 + \frac{1+\delta}{\delta} \frac{\log \log N}{\log N} + \frac{\log x}{\log N} \right) \right\rceil.$$

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 urn defined by Equation (9) is such that the variable

$$\frac{(\log N)^{(1+\delta)/\delta}}{(\alpha\delta N)^{1/\delta}} \nu^D(N) - \log N - \frac{1+\delta}{\delta} \log \log N \quad (19)$$

converges in distribution to a random variable Y defined by

$$\mathbb{P}(Y \geq x) = \exp\left(-(\alpha\delta)^{1/\delta} e^x\right), \quad x \in \mathbb{R}.$$

PROOF. Chen-Stein's method is the basic tool in the proof of the theorem. See Barbour *et al.* [4] 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 urns when the balls in the i_0 -th urn are thrown again until the i_0 -th urn is empty. It follows that the number of

balls in any other urn 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.* [4]. Then, by [4, Corollary 2.C.2], the following relation holds,

$$\begin{aligned} \sum_{p \geq 0} \left| \mathbb{P}(W_n^k = p) - \frac{\mathbb{E}(W_N^k)^p}{p!} e^{-\mathbb{E}(W_N^k)} \right| \\ \leq 1 - \text{Var}\left(W_N^k\right) / \mathbb{E}\left(W_N^k\right). \end{aligned}$$

By taking $k = \kappa_x(N)$ and by using Propositions 3 and 4, 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 $\mathbb{P}(W_N^k = 0) = \mathbb{P}(\nu^D(N) > k)$. \square

The convergence in distribution of $\nu^D(N)$ has been proved by Csáki and Földes [6] with a different method. Our result gives a more accurate description of the location of empty urns (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 1 (CUTOFF PHENOMENON). *Under the assumption of Theorem 1, if*

$$k(N) = (N/\log N)^{1/\delta},$$

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 urns have been considered. The result below shows that the first empty urn happens at a time of the order of $\log N$. Remembering the approximation of the peer to peer system, it suggests that the time the system begins to serve quickly the incoming peers should be of the same order.

COROLLARY 2 (FIRST EMPTY URN). *Let*

$$T^D(N) = T_{\nu^D(N)} = \sum_{k=1}^{\nu^D(N)} \frac{E_k^1}{k}. \quad (20)$$

Under the assumptions of Theorem 1, the quantity

$$\delta T^D(N) - \log N + \log \log N - \log(\alpha\delta)$$

converges in distribution to T_∞ , where T_∞ is the random variable defined in Proposition 2.

PROOF. If V_N is the variable defined by Expression (19), then

$$\begin{aligned} \delta \log \nu^D(N) - \log(N) + \log \log N - \log(\alpha\delta) \\ = \delta \log \left(\frac{1}{\log N} \left(V_N + \log N + \frac{1+\delta}{\delta} \log \log N \right) \right). \end{aligned}$$

Since by Theorem 1 the sequence (V_n) converges in distribution, it implies that the right hand side of the above expression converges in distribution to 0. Proposition 2 shows that

$$E_1^1 + \frac{E_2^1}{2} + \dots + \frac{E_n^1}{n} - \log n$$

converges almost surely to T_∞ . \square

5. RANDOM PROBLEM

For the random model, the probability P_n of selecting the n -th urn is given by Equation (4) of Proposition 2. In the (almost sure) limit as n goes to infinity, $X_n \sim X_\infty$ 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 urn. The fact that Z_n , related to the “width” of the n -th urn, can be arbitrarily small with a positive probability suggests that the index ν^R of the first empty urn 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 [12], and only sketches of proof are given for Proposition 5 and Theorem 2 in the present paper. It must be noticed that a similar problem where X_∞ and the sequence $(E_n^1, n \geq 1)$ are independent is fairly easy to solve. However here, these random variables are dependent, and this dependency requires quite technical probabilistic tools.

To derive asymptotic results for ν^R , as in the previous section, the asymptotic behavior of the random variable W_N^k defined by

$$W_N^k = \sum_{i=1}^k I_{N,i} \text{ with } I_{N,i} = \mathbb{1}_{\{\eta_i^R(N)=0\}}.$$

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 urns 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), which suggests that if a limit in distribution exists, it cannot be Poisson.

As was pointed out in Hwang and Janson [10], the sequence $(NP_i, 1 \leq i \leq k)$ plays a central role in the limiting behavior of (W_N^k) . 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 $\nu^R(N)$; see [12] for the proof.

PROPOSITION 5. *Let $x > 0$. When N goes to infinity, the random sequence $(NP_i, 1 \leq i \leq xN^{1/(\rho+2)})$ converges in distribution to a doubly stochastic Poisson process with a random intensity $x^{\rho+2}(X_\infty \rho(\rho+2))^{-1}$.*

PROOF. Because of the technicality involved, we only give a sketch of the proof. The reader is referred to [12] for more details.

To prove the convergence of the sequence of point processes $\mathcal{N}_N = \sum_{i=1}^{k(N)} \delta_{\{NP_i\}}$ with $k(N) = xN^{1/(\rho+2)}$, it is enough to show the convergence of the Laplace transforms of these point processes applied to some suitable functions. Non-negative continuous functions with a compact support would be enough to prove the result, but the next theorem requires a slightly stronger result, namely it requires the convergence of Laplace transforms for non-negative continuous functions vanishing at infinity, i.e., that for any function $f \geq 0$ continuous vanishing at infinity, we have

$$\lim_{N \rightarrow +\infty} \mathbb{E} \left(e^{-\mathcal{N}_N(f)} \right) = \mathbb{E} \left(e^{-\mathcal{N}_\infty(f)} \right)$$

where, conditionally on X_∞ , \mathcal{N}_∞ is a Poisson process with intensity $x^{\rho+2} X_\infty^{-1} / (\rho+2)$.

The general idea is to condition on the random variable X_∞ . However, for each $n \geq 1$, X_∞ and Z_n are dependent, so that this cannot be directly done. Instead, the first step of the proof is to show that only the last terms of the point process matter, i.e., that $\mathbb{E}(e^{-\mathcal{N}_N(f)})$ and $\mathbb{E}(e^{-\sum_{\beta(N)}^{k(N)} f(NP_i)})$ have the same limit, for any sequence $\beta(N) \ll k(N)$. So we are left with large indexes $i \geq \beta(N)$, for which the approximation $P_i = \rho i^{-\rho-1} X_i Z_i \approx \rho i^{-\rho-1} X_{\beta(N)} Z_i$ can be justified. The main tool behind this approximation is Doob’s Inequality applied to the reversed martingale

$$\overline{M}_n = \sum_{k \geq n} (E_k - 1) / k.$$

And now, due to this approximation, it is perfectly rigorous to condition on $\mathcal{F}_N = \sigma(E_k, k < \beta(N))$: since for $i \geq \beta(N)$, Z_i is independent of $X_{\beta(N)}$, we are exactly left with proving the result for the sequence of point processes $\mathcal{N}'_N = \sum_{\beta(N)}^{k(N)} \delta_{\{N x_N i^{-\rho-1} Z_i\}}$ with any converging sequence $x_N \rightarrow x_\infty$ (x_N has to be thought as being equal to $\rho X_{\beta(N)}$). If f has a compact support, it is possible to conclude by applying a result from Grigelionis [9] to show that this sequence of point processes converges to a Poisson process with intensity $x^{\rho+2} / (x_\infty(\rho+2))$. In the general case, the convergence is shown thanks to computations, by controlling the speed at which Z_i converge in law to an exponential random variable. \square

This result together with standard poissonization techniques make it possible to prove the following theorem, which is the main result of this section.

THEOREM 2. *Let $\kappa(N) = N^{1/(\rho+2)}$. For $x > 0$, $W_N^{x\kappa(N)}$ converges in distribution to a Poisson random variable with a random parameter $x^{\rho+2}(X_\infty \rho(\rho+2))^{-1}$ when $N \rightarrow \infty$.*

PROOF. Again, only a sketch of the proof is given. The first step of the proof is to show the result for the random variable $W_{\mathcal{P}_N}^{x\kappa(N)}$ where \mathcal{P}_N is a Poisson random variable with parameter N , independent of everything else so far. The idea is that the law of $W_{\mathcal{P}_N}^{x\kappa(N)}$ is not sensitive to the fluctuations of \mathcal{P}_N around its mean value, equal to N , so that the law of $W_{\mathcal{P}_N}^{x\kappa(N)}$ and of $W_N^{x\kappa(N)}$ will have the same asymptotic behavior.

To show the convergence of $W_{\mathcal{P}_N}^{x\kappa(N)}$, we consider its generating function: for $u > 0$ and $k \in \mathbb{N}$, we can compute

$$\mathbb{E}\left(u^{W_{\mathcal{P}_N}^k}\right) = \mathbb{E}\left(e^{\sum_{i=1}^k \log(1-(1-u)e^{-NP_i})}\right) = \mathbb{E}\left(e^{-\mathcal{N}_{N,k}(fu)}\right),$$

where $\mathcal{N}_{N,k} = \sum_{i=1}^k \delta_{\{NP_i\}}$, and $f_u(x) = -\log(1 - (1-u)e^{-x})$ for $x \geq 0$. Then $\int_0^\infty (1-e^{-f_u}) = 1-u$, so that we conclude with the previous proposition that $W_N^{x\kappa(N)}$ converges to a random variable which, conditionally on X_∞ , is a Poisson random variable with parameter $x^{\rho+2}(X_\infty \rho(\rho+2))^{-1}$. The fact that $W_N^{x\kappa(N)}$ and $W_{\mathcal{P}_N}^{x\kappa(N)}$ have the same asymptotic behavior (in law) then follows by standard arguments. \square

This theorem readily yields the following corollary.

COROLLARY 3. *The random variable $\nu^R(N)/\kappa(N)$ converges in distribution to a random variable Y such that*

$$\mathbb{P}(Y \geq x) = \mathbb{E}\left(e^{-x^{\rho+2}X_\infty^{-1}/(\rho(\rho+2))}\right).$$

Finally, if $T^R(N) \stackrel{\text{def}}{=} T_{\nu^R(N)}$ then, for the convergence in distribution,

$$\lim_{N \rightarrow +\infty} \frac{T^R(N)}{\log(N)} = \frac{1}{\rho+2}. \quad (21)$$

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 2 that $\lim \mathbb{E}(W_N^{x\kappa(N)})$ is proportional to $\mathbb{E}X_\infty^{-1}$ and $\mathbb{E}X_\infty^{-1} < +\infty$ 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 $W_N^{x\kappa(N)}$ diverges because it happens that a finite number of intervals (actually, the $\lfloor \rho \rfloor$ first intervals) capture most of the balls. This 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 paper in term of urn and ball models. These simulations are in particular used to justify Approximation 1, as well as to compare the insights into the dynamics of the

system provided by the two urn and ball models studied in this paper. Moreover, another server selection policy is considered, namely when an incoming peer chooses the server at random.

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, we define different times:

1. \tilde{T}_1 is the first time when two servers are created and less than 2 peers have arrived.
2. \tilde{T}_2 is the last time when there is an empty server.
3. \tilde{T}_3 is the first time when the input rate is smaller than the output rate (see Section 6.3).
4. \tilde{T}_4 the first time when a server becomes empty, i.e., when a peer leaves a server where it was alone.

We consider the corresponding quantities $\tilde{\nu}_i$: for $i = 1, 2, 3, 4$, $\tilde{\nu}_i$ is the index of the interval (S_{i-1}, S_i) in which the event corresponding to \tilde{T}_i happens. In particular, $\tilde{\nu}_1$ corresponds to Definition 1. In every simulation, the averages of the quantities $\tilde{\nu}_i$ and \tilde{T}_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 $5 \cdot 10^7$.

6.1 Validation of Approximation 1

Definition 1 specifies the variable considered throughout this paper to determine the duration of the first regime of the file sharing system. This variable was chosen for two reasons. First, it is a good indicator of the current equilibrium of the system: the output rate begins to be comparable with the input rate when only a few peers arrive between the creation of two successive servers. Moreover, the stopping time defined in this way is mathematically tractable when transposed into the context of a certain urn and ball model. Compared to [17], it is interesting to note that we are actually able to rigorously prove results, and not only rely on simulation. As a byproduct, the mathematical problems arising in this context are interesting in themselves.

For the sake of completeness, several points need to be addressed. First, for how long is Approximation 1 valid? Since the random variable ν specified in Definition 1 corresponds to $\tilde{\nu}_1$, we argued in Section 2.3 that this approximation holds until \tilde{T}_1 . This is the main assumption that makes it possible to cast our problem in terms of urns and balls, and to derive precise results on $\tilde{\nu}_1$ and \tilde{T}_1 .

In order to validate the results of Section 5, we check that $\tilde{\mathbb{E}}(\nu_1)$ and $\mathbb{E}(\tilde{T}_1)$ behave as predicted by Theorem 2. From this theorem, we expect to have $\mathbb{E}(\tilde{\nu}_1) \approx A_1 N^{1/(\rho+2)}$ for some constant A_1 , and $\mathbb{E}(\tilde{T}_1) \approx \log(N)/(\rho+2)$. Figure 2 shows the graphs $\log(\mathbb{E}(\tilde{\nu}_1))$ and $\mathbb{E}(\tilde{T}_1)$ versus $\log(N)$: the straight lines depicted prove a good agreement with the theory. Moreover, via a fitting procedure, one can compute the slopes of these lines: the results are summarized in Table 1.




Figure 2: $\log(\mathbb{E}(\tilde{\nu}_1))$ (solid) and $\mathbb{E}(\tilde{T}_1)$ (dashed), $\rho = 2$.

The values of interest in Table 1 are in the row labeled “Min”: we see that simulations exhibit a slope of 0.248 for $\tilde{\nu}_1$ and of 0.256 for \tilde{T}_1 , whereas the theory predicts 0.25 in both cases (because $\rho = 2$). These results are in good agreement with Approximation 1, which justifies the fact that we can use this approximation up to time \tilde{T}_1 .

Table 1: Coefficients of growth rates of Fig. 2 and 3 in the case $\rho = 2$

Policy	$\tilde{\nu}_1$	\tilde{T}_1	$\tilde{\nu}_2$	\tilde{T}_2	$\tilde{\nu}_4$	\tilde{T}_4
Min	0.2478	0.2565	0.3765	0.5146	0.3149	0.3287
Random	0.2470	0.2575	0.3711	0.5078	0.2383	0.2530

6.2 Accuracy of Urn and Ball Models

In this section, we compare the random and deterministic urn and ball models with $\mathbb{E}(\tilde{T}_2)$, the expected value of the last time when there is an empty server. It clearly appears in Figure 1 that \tilde{T}_2 closely corresponds to the shift in equilibrium of the system, and this fact has been observed in numerous simulations. However, as we will see in the following, Approximation 1 does not hold until time \tilde{T}_2 , which explains why it is very challenging from a mathematical point of view to derive results on T_2 . (Note in addition that \tilde{T}_2 is not a stopping time).

Figure 3 shows that \tilde{T}_1 is much smaller than \tilde{T}_2 : This result is nevertheless not surprising. Indeed, as discussed in Section 2, results obtained for the random model point out a local behavior: the first empty urn arrives in a region, where still many peers arrive in each interval. Although many peers should arrive in this time interval, this is in reality not the case because a very small interval is generated. Thus, in some sense, the order of magnitude $N^{1/(\rho+2)}$ provided by the random urn and ball model is misleading for the initial system.

In the deterministic model, the sizes of urns are not random,




Figure 3: The times $\mathbb{E}(\tilde{T}_1) \leq \mathbb{E}(\tilde{T}_2)$ and $\mathbb{E}(T^D)$ when $\rho = 2$.

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 defined by Equation (20) predicted by the deterministic model is closer to \tilde{T}_2 than to \tilde{T}_1 .

Although considering the deterministic model indeed improves the approximation, \tilde{T}_2 still seems much larger than T^D . However, thanks to our urn and ball models, we know that the first order approximation of the times \tilde{T}_i is logarithmic, whereas the first order approximation for the indexes $\tilde{\nu}_i$ is polynomial. Table 1 provides useful information to understand the situation.

First, the deterministic model yields a reasonable estimate of the exponent in $\tilde{\nu}_2$: simulations give 0.376 and the deterministic model 0.333. Note that the random model predicts 0.25, so a substantial improvement in accuracy is obtained when using the deterministic model. This suggests that Approximation 1 holds until T^D , i.e., up to times of order $N^{1/(\rho+1)}$.

Second, we observe a significant discrepancy between the exponent for $\tilde{\nu}_2$ and the coefficient of \tilde{T}_2 : If Approximation 1 were to hold until \tilde{T}_2 , one would have $\tilde{T}_2 \approx \sum_{k=1}^{\tilde{\nu}_2} E_k^1/k$, which would yield, because $\tilde{\nu}_2 \approx N^{0.38}$, that $\tilde{T}_2 \approx 0.38 \log(N)$. However, we find that the time \tilde{T}_2 is better approximated by $0.52 \log(N)$, and so Approximation 1 does not hold until time \tilde{T}_2 . This clearly poses the challenge to derive asymptotic results for \tilde{T}_2 . Moreover, this triggers another interesting question: For how long does Approximation 1 hold? We give a partial answer to this question by considering the times \tilde{T}_3 and \tilde{T}_4 in the next section.

6.3 On the Duration of Approximation 1

Throughout this paper, we have tried to estimate the time when the equilibrium of the system begins to shift. As long as Approximation 1 holds, the input rate $i(t)$ of the system

is the number of peers, that are not active at time t , times ρ , while the output rate $o(t)$ is just the number of servers at time t (since the service has mean one). Initially, $i(0) = \rho N$ and $o(0) = 1$, and $i(\infty) = 0$ and $o(\infty) = N$. To study the time at which the equilibrium of the system begins to shift, it is therefore very natural to consider the first time \tilde{T}_3 at which $i(t) < o(t)$, i.e., when the number of servers is greater than ρ times the number of non-active peers. As shown in the following, considering this time leads to the order of magnitude given by the deterministic model (with less precise asymptotics of course).

For times $t < \tilde{T}_3$, we assume that Approximation 1 holds, so that we can cast $\tilde{\nu}_3$ in terms of our urn and ball problem. Let Z_N^x be the number of balls that fall in the x first intervals:

$$Z_N^x = \sum_{i=1}^x \eta_i(N) = \sum_{i=1}^N \mathbb{1}_{\{E_i^\rho \leq T_x\}}.$$

The index ν_3 then corresponds to

$$\nu_3 = \inf \left\{ x : N - Z_N^x \stackrel{\text{def}}{=} \tilde{Z}_N^x < \frac{x}{\rho} \right\}.$$

The asymptotic behavior of $\mathbb{E}(\tilde{Z}_N^x)$ when x goes to infinity with N is easy to derive:

$$\mathbb{E}(\tilde{Z}_N^x) = N \sum_{i>x+1} \mathbb{E}P_i \sim \alpha N \sum_{i>x+1} i^{-\rho-1} \sim \frac{\alpha}{\rho} N x^{-\rho}$$

Therefore $\mathbb{E}(\tilde{Z}_N^x) \approx x$ for $x \approx N^{1/(\rho+1)}$, i.e. $\tilde{\nu}_3$ is of order $N^{1/(\rho+1)}$, which is the same order of magnitude as in the deterministic model. Rigorous mathematical analysis could be done to prove this result, but in our view, considering \tilde{T}_1 has one main advantage: Proposition 1 is almost a rigorous justification of Approximation 1. When considering another time, in particular \tilde{T}_3 , we were not able to provide such a strong justification. And as we have seen in the case of \tilde{T}_2 , Approximation 1 does not hold for the whole first regime, and a strong justification as Proposition 1 is therefore very valuable.

Finally, let us give some brief results on \tilde{T}_4 , the first time when a server empties. Simulations show that $\tilde{\nu}_4$ and \tilde{T}_4 have similar behavior as before (polynomial and logarithmic growths, respectively). Results in Table 1 show that the slope for \tilde{T}_4 is similar to the exponent of $\tilde{\nu}_4$, suggesting that Approximation 1 holds until \tilde{T}_4 .

In conclusion, Approximation 1 holds at least until $N^{1/(\rho+1)}$, which corresponds to \tilde{T}_1 and \tilde{T}_3 . However, it does not hold until \tilde{T}_2 , whereas Figure 1 shows that until \tilde{T}_2 , the system is still in the first regime. For the particular value $\rho = 2$, we have $\nu^D \approx AN^{0.33}$ and simulations show that $\tilde{\nu}_2 \approx A_2 N^{0.38}$, and so our approximation by the means of a urn and ball problem is not so far from the exponent that we want to capture. Proposition 1 shows that until ν^D , there are only few empty servers: so between T^D and \tilde{T}_2 , it could happen that there is a fraction of empty servers, and although this fraction is small, it has an impact on the system. Similar phenomenon have been observed in Sanghavi *et al.* [15].

To conclude this section, we discuss a different routing policy. Throughout this paper, we have considered the policy

where an incoming peer selects the least loaded server, in terms of number of peers. This policy is compared against the random one, where an incoming peer selects a server uniformly at random among all possible servers.

Simulations show that these policies are very close as shown in Figures 4, 5 and 6. The only noticeable difference is concerning $\mathbb{E}(\tilde{\nu}_3)$, cf. Figure 7. However, Table 1 shows that the exponents of $\mathbb{E}(\tilde{\nu}_3)$ are very similar in the random and in the minimum policy. One can easily check that they are indeed proportional one to another.




Figure 4: Comparison of Min and Random for $\mathbb{E}(\tilde{T}_1)$ when $\rho = 2$

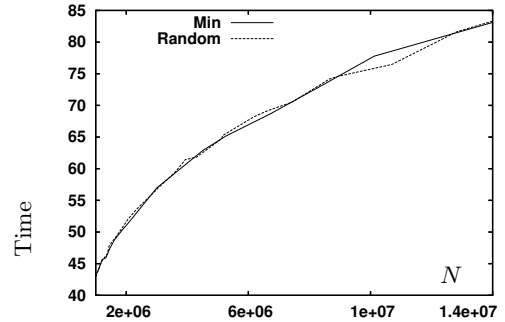


Figure 5: Comparison of Min and Random for $\mathbb{E}(\tilde{\nu}_1)$ when $\rho = 2$

Table 1 shows that for the first time when a server becomes empty, the policy has a great influence. This is easily understandable: In the min case, it is much harder for a server to become empty, because least loaded servers are selected by incoming peers.

7. CONCLUSION

The simulations moreover underlined the existence of a second regime during which although the fraction of idle servers is small, the output rate is no longer as high as possible. This second regime is then followed by a third regime during which the capacity offered by the system exceeds by far the input rate, and so the system mainly creates empty servers. Our urn and ball approach can no longer be applied to these two regimes, and so they will be studied in the near future using other probabilistic techniques.

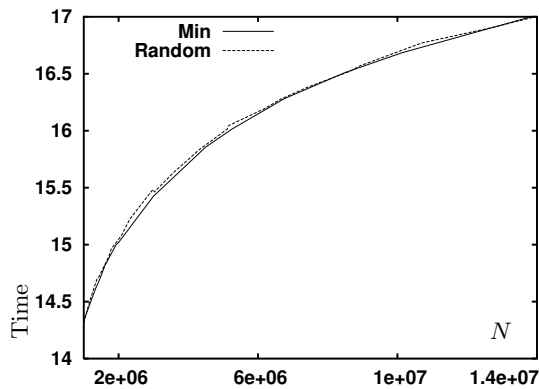


Figure 6: Comparison of Min and Random for $\mathbb{E}(\tilde{T}_2)$ when $\rho = 2$

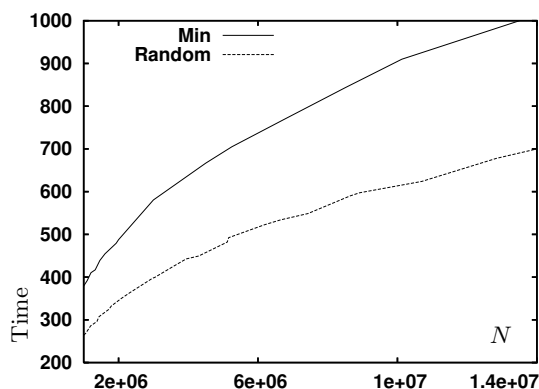


Figure 7: Comparison of Min and Random for $\mathbb{E}(\tilde{v}_2)$ when $\rho = 2$

A possible extension of our results consists of incorporating the possibility for a peer to leave the system right after completing its download. In terms of urn and ball, this just amounts to change the parameter that defines the length of the n -th interval: instead of n , one would just have to consider pn if p is the probability for a peer to become a server after completing its download. An extended model where the file is split into different chunks essentially amounts to study a multi-class queueing network with a random number of servers of different classes which proves to be a much more difficult problem.

Finally, a natural extension is to consider a general service distribution, instead of the exponential one. In this case, the process of creation of servers can be described as an age-dependent branching process, and more precisely a binary Bellman-Harris branching process. See Athreya [2, 3]. If this setting complicates significantly the analysis of the file sharing system, it seems that most of the results obtained in the exponential case should still hold.

8. REFERENCES

- [1] Søren Asmussen, *Applied probability and queues*, John Wiley & Sons Ltd., Chichester, 1987.
- [2] K. B. Athreya and Niels Keiding, *Estimation theory for continuous-time branching processes*, Sankhya: The Indian Journal of Statistics **89** (1977), no. A, 101–123.
- [3] Krishna B. Athreya, *On the supercritical one dimensional age dependent branching processes*, The Annals of Mathematical Statistics **40** (1969), no. 3, 743–763.
- [4] A. D. Barbour, Lars Holst, and Svante Janson, *Poisson approximation*, The Clarendon Press Oxford University Press, New York, 1992, Oxford Science Publications.
- [5] F. Clévenot and P. Nain, *A simple fluid model for the analysis of the Squirrel peer-to-peer caching system*, Proc. Infocom, 2004.
- [6] E. Csáki and A. Földes, *On the first empty cell*, Studia Scientiarum Mathematicarum Hungarica **11** (1976), no. 3-4, 373–382 (1978).
- [7] Z. Ge, D.R. Figueiredo, S. Jaiswal, J. Kurose, and D. Towsley, *Modeling peer-to-peer file sharing systems*, Proc. Infocom, 2003.
- [8] Alexander Gnedin, Ben Hansen, and Jim Pitman, *Notes on the occupancy problem with infinitely many boxes: general asymptotics and power laws*, Probability Surveys **4** (2007), 146–171 (electronic). MR MR2318403
- [9] Bronius Grigelionis, *On the convergence of sums of random step processes to a poisson process*, Theory of Probability and its Applications **8** (1961), no. 2, 177–182.
- [10] Hsien-Kuei Hwang and Svante Janson, *Local limit theorems for finite and infinite urn models*, Annals of Probability (To appear).
- [11] J. F. C. Kingman, *Random partitions in population genetics*, Proceedings of the Royal Society. London. Series A. Mathematical, Physical and Engineering Sciences **361** (1978), no. 1704, 1–20.
- [12] Unpublished manuscript available at, <http://chambertin.inria.fr/robert/Sigmetrics-Theorem2.pdf>.
- [13] L. Massoulié and M. Vojnović, *Coupon replication systems*, Proc. Sigmetrics 2005 (Banff, Alberta, Canada), June 2005.
- [14] D. Qiu and R. Srikant, *Modeling and performance analysis of BitTorrent-like peer-to-peer networks*, Proc. Sigcomm, 2004.
- [15] Sujay Sanghavi, Bruce Hajek, and Laurent Massoulié, *Gossiping with Multiple Messages*, INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE, May 2007, pp. 2135–2143.
- [16] David Williams, *Probability with martingales*, Cambridge University Press, 1991.
- [17] Xiangying Yang and Gustavo de Veciana, *Performance of peer-to-peer networks: service capacity and role of resource sharing policies*, Performance Evaluation **63** (2006), no. 3, 175–194.