

A SCALING ANALYSIS OF A TRANSIENT STOCHASTIC NETWORK

MATHIEU FEUILLET * ** AND

PHILIPPE ROBERT, * *** INRIA-Rocquencourt

Abstract

In this paper we use a simple transient Markov process with an absorbing point to investigate the qualitative behavior of a large-scale storage network of nonreliable file servers across which files can be duplicated. When the size of the system goes to ∞ , we show that there is a critical value for the maximum number of files per server such that, below this quantity, most files have a maximum number of copies. Above this value, the network loses a significant number of files until some equilibrium is reached. When the network is stable, we show that, with convenient time scales, the evolution of the network towards the absorbing state can be described via a stochastic averaging principle.

Keywords: Distributed system with breakdown; time scale; transient Markov chain with absorbing state; Skorokhod problem

2010 Mathematics Subject Classification: Primary 60K25; 60F05

Secondary 68M14; 90B05

1. Introduction

1.1. Storage systems

We consider a large-scale storage system, namely a set of file servers in a communication network. In order to ensure persistence, files are duplicated on several servers. When the disk of a given server breaks down, its files are lost but can be retrieved on the other servers if copies are available. For this reason, a fraction of the bandwidth of a server is devoted to the duplication mechanism of its files to other servers. On the one hand, there should be sufficiently many copies so that any file has a copy available on at least one server at any time. On the other hand, in order to use the bandwidth in an optimal way, there should not be too many copies of a given file so that the network can accommodate a large number of distinct files. These systems are known as distributed hash tables (DHTs). They play an important role in the development of some large-scale distributed systems; see [30] and [32] for a more detailed presentation.

Failures of disks occur naturally randomly; these events are quite rare, but, given the large number of nodes in these distributed systems, this is not a negligible phenomenon at the level of the network. If, for a short period of time, several of the servers break down, it may happen that files will be lost for good because all the available copies were on these servers and because the recovery procedure was not completed before the last copy disappeared. When designing such a system, it is therefore desirable to find a convenient duplication policy and to design the system so that a copy of each file exists for as long as possible. The natural critical parameters of the network are the failure rates of the servers, the bandwidth allocated to duplication, the

Received 7 December 2012; revision received 22 April 2013.

* Postal address: INRIA-Rocquencourt, Domaine de Voluceau, 78153 Le Chesnay, France.

** Email address: mathieu.feuillet@inria.fr

*** Email address: philippe.robert@inria.fr

number of files, and the number of servers. The ratio of the two last quantities is a measure of the storage capacity of the system. It is important to understand the impact of each of these parameters on the efficiency of the storage system.

1.2. Literature

This network can be seen as a classical set of queues with breakdowns. Numerous stochastic models of such systems have been investigated in the literature; see Chapter 6 of [21] for example and the references therein. Related models include queues with retrial and queues with servers of walking types; see [2] and [10]. For most of the systems analyzed, there are, in general, one or two nodes that are subject to breakdown. A queueing analysis is generally done in this context: convergence in distribution of the associated Markov model and analysis of the distributions of the availability of the system, of the delays, and of the queue sizes. There are very few stochastic models that can be used to investigate the performance of DHTs; the existing models describe the evolution of the number of copies of a single file. See [5], [27], and [28]; see also [13]. In most of these studies the interaction between different files, due to the bandwidth sharing limitations, has not really been considered, except through simulations. The purpose of this paper is to investigate the impact of this interaction. We consider a large system, i.e. a large number of files, instead of a small number of elements, which is a natural consideration for current distributed systems.

1.3. The stochastic model

We consider the following simple model. A file can have at most two copies, with the total bandwidth allocated to file duplication given by λN for $\lambda > 0$ and $N \in \mathbb{N}$. If at some moment there are $x \geq 1$ files with exactly one copy, a new copy of each of these files is created at rate $\lambda N/x$. It is assumed that initially F_N files are present in the system with two copies and each copy of a file disappears at rate μ . Recall that a file with 0 copies is lost. It will be assumed that the total number of files F_N is proportional to N , i.e. that F_N/N converges to some $\beta > 0$. Clearly, this system is transient and the empty state, i.e. all files are lost, is an absorbing state. The aim of this paper is to describe the decay of the network, i.e. how the set of lost files increases. For $\delta > 0$, there exists some finite random instant $T_N(\delta)$ such that $\lfloor \delta N \rfloor$ files are lost after time $T_N(\delta)$. In this paper we investigate the order of magnitude in N of the variables $T_N(\delta)$ as N becomes large, and the role of the parameters λ , μ , and β in these asymptotics.

1.4. Modeling issues

As will be seen, one of the conclusions of this analysis is that if the parameters λ , μ , and β of the network are properly dimensioned, then the system remains for a long period of time in a kind of stationary state where only a few files are permanently lost. Curiously, this important property does not seem to have been explicitly investigated by the designers of algorithms for these storage systems, even if it was probably suspected. In general, algorithms are evaluated experimentally with an overloaded initial state so that a significant fraction of files is lost at the start. The ‘best’ algorithm is that which loses files less rapidly; see [23] for example. These experimental simulations are conducted for a limited period of time, e.g. several weeks, during which there is a sharp degradation, indicating that the decay to the empty state follows. However, as will be seen with a simple model, this is misleading; there is a stabilization of the network and this is why it makes sense to design such systems.

In order to concentrate on these stability properties, we neglect some other important issues concerning these storage systems. We review some of them.

Fixed set of files. In practice, new files arrive that have to be stored and duplicated. Initially, a large number of files need to be stored and we study the conditions under which this initial set of files can be preserved for some period of time. In the case of a Markovian model with arrivals, one could consider the associated invariant distribution. This has been done, partially, in the literature. In our opinion such a study would perhaps miss the interesting transient properties which are exhibited here.

Capacity for duplication. If there are N servers, each with an available bandwidth λ to duplicate files, then the maximal capacity for duplication is λN . The model described above has therefore an optimal use of the duplication mechanism since the maximal duplication capacity is always available. For this reason, this model provides upper bounds on the optimal performance of such a system. In particular, for any duplication mechanism, after a duration of time with the same distribution as $T_N(\delta)$, at least $\lfloor \delta N \rfloor$ files will be lost for good.

Homogeneity. In our study all files have the same loss rate, which is unlikely in practice since the storage servers may have different reliability characteristics. To extend our results, we could describe the state of the systems by adding a variable to indicate the loss rate of a given copy. However, we work in a homogeneous framework for simplicity.

1.5. Time scales of transient Markov processes

If, for $i \in \{0, 1, 2\}$, $X_i^N(t)$ denotes the number of files with i copies in the network, then, under Poisson assumptions for failures and duplication processes, $(X_0^N(t), X_1^N(t))$ is clearly a finite Markov process with $(F_N, 0)$ as an absorbing state. Contrary to the abovementioned works, there is clearly no question of equilibrium here since the system *dies* at $(F_N, 0)$. A possible approach to investigate the decay of such a system is to consider the associated quasistationary distributions of the Markov process; see [6] and [11] for example. It would give a description of the system conditionally on the event that only a fraction of the files has been lost. It is in general expressed in terms of the spectral characteristics of the jump matrix. For this reason, an explicit description of these distributions is quite rare outside one-dimensional birth-and-death processes. In this paper, different time scales will be used to investigate the qualitative behavior of these transient processes. Times scales can be thought of as ‘lenses’.

We show that a stochastic averaging principle (SAP) occurs for this transient process; roughly speaking, its dynamics depend on two components, one evolving over a fast time scale and the other evolving over a slower time scale. See [25, Chapters 9 and 10] for an elementary introduction, and also [17] and [14]. This phenomenon is known to occur for the classical example of loss networks. In this case the vector of the number of free places of the congested links is the fast component; see [20] and [18]. Outside this class of network, there are currently few examples of stochastic networks for which such a SAP occurs. See [12] and [26] for recent examples of SAPs. In our case a SAP occurs as follows: under the stability condition $\lambda > 2\beta\nu$, the transition rates associated with exactly one copy are proportional to N so that it is instantaneously at its equilibrium depending on the current fraction of lost files. The number of lost files increases at a bounded rate; hence, the process evolves on a much slower time scale. This intuitive phenomenon requires some technical steps to be established rigorously. In this paper, a martingale approach is used to tackle the technical intricacies. See Theorem 4 and the discussion following its proof.

1.6. Outline of the paper

In Section 2 we introduce the Markov process investigated and its corresponding martingale representation.

In Section 3 we study a fluid picture of the network, i.e. the limit of the sequence of processes $(X_0^N(t)/N, X_1^N(t)/N)$. We show in Theorem 1 that its limit, the solution of an ordinary differential equation, is not trivial when $\lambda < 2\mu\beta$ and is $(0, 0)$ when $\lambda > 2\mu\beta$. The storage system is therefore properly designed when $\lambda > 2\mu\beta$; otherwise, it is inefficient since it loses a significant number of files from the start.

Section 4 is devoted to the critical case $\lambda = 2\mu\beta$. In Theorem 2 we show that the sequence of processes $(X_0^N(t)/\sqrt{N}, X_1^N(t)/\sqrt{N})$ converges in distribution and that its limit can be expressed in terms of a non-Markovian one-dimensional process, which is the solution of an unusual stochastic differential equation with reflection at 0.

In Section 5 we investigate the stable case $\lambda > 2\mu\beta$. We show that the capacity of the system remains intact at the normal time scale: for $t \geq 0$, in Theorem 3 we prove that the variable $(X_0^N(t))$ converges in distribution to a Poisson process. Only a finite number of files is lost as N goes to ∞ .

The main result of the paper, Theorem 4 describes on the time scale $t \rightarrow Nt$ the transience of the Markov process: a fraction $\psi(t)N$ of the files is lost, where $\psi(t)$ is the solution of some fixed-point equation. Consequently, $t \rightarrow Nt$ is a convenient time scale to use to observe the degradation of the storage system. In this case a stochastic averaging principle holds: around time Nt there is a local equilibrium for which $(\beta - \psi(t))N$ files are still available.

The proof of the convergence results uses an extension of the classical Skorokhod problem proposed in [1]. The necessary material is gathered in Appendix A to keep the paper self-contained.

2. The stochastic model

Recall that F_N is the total number of distinct files initially present in the network, and that $X_1^N(t)$ and $X_0^N(t)$ are respectively the number of files with one copy at time t and the number of lost files at time t . The number $X_2^N(t)$ of files with two copies at time t is defined by $X_2^N(t) = F_N - X_0^N(t) - X_1^N(t)$.

We assume the initial state

$$(X_0^N(0), X_1^N(0)) = (0, 0),$$

and that all files initially have two copies. The copy of a file is lost with rate μ and, conditionally on $X_1^N(t) = x$, a file with only one copy gets an additional copy with rate $\lambda N/x$. All events are supposed to occur after an exponentially distributed time. Under these assumptions, $(X(t)) := (X_0^N(t), X_1^N(t))$ is a Markov process on the state space

$$\mathcal{S} = \{x = (x_0, x_1) \in \mathbb{N}^2 : x_0 + x_1 \leq F_N\};$$

as mentioned above, under these assumptions, the state $(F_N, 0)$ is an absorbing point of the process $(X^N(t))$.

For $x \in \mathbb{N}^2$, the Q -matrix $Q^N = (q^N(\cdot, \cdot))$ of the process $(X(t))$ is defined by

$$q^N(x, x + e_1) = 2\mu(F_N - x_0 - x_1), \tag{1a}$$

$$q^N(x, x - e_1) = \lambda N \mathbf{1}_{\{x_1 > 0\}}, \tag{1b}$$

$$q^N(x, x - e_1 + e_0) = \mu x_1, \tag{1c}$$

with $e_0 = (1, 0)$ and $e_1 = (0, 1)$. We assume that

$$\lim_{N \rightarrow +\infty} \frac{F_N}{N} = \beta,$$

and define $\rho = \lambda/\mu$.

The stochastic differential equations associated with this transient Markov process can be written as

$$X_0^N(t) = X_0^N(0) + \sum_{i=1}^{F_N} \int_0^t \mathbf{1}_{\{i \leq X_1^N(u-)\}} \mathcal{N}_{\mu,i}(du), \tag{2}$$

$$\begin{aligned} X_1^N(t) = & X_1^N(0) - \int_0^t \mathbf{1}_{\{X_1^N(u-) > 0\}} \mathcal{N}_{\lambda N}(du) - \sum_{i=1}^{F_N} \int_0^t \mathbf{1}_{\{i \leq X_1^N(u-)\}} \mathcal{N}_{\mu,i}(du) \\ & + \sum_{i=1}^{F_N} \int_0^t \mathbf{1}_{\{i \leq F_N - X_0^N(u-) - X_1^N(u-)\}} \mathcal{N}_{2\mu,i}(du), \end{aligned} \tag{3}$$

where $(\mathcal{N}_{\mu,i})$ and $(\mathcal{N}_{2\mu,i})$ are two independent sequences of independent and identically distributed Poisson processes with respective parameters μ and 2μ , and $\mathcal{N}_{\lambda N}$ is an independent Poisson process with parameter λN . For the i th file having only one copy, the integrand on the right-hand side of (2) corresponds to its definitive loss and the second term on the right-hand side of (3) is associated with its duplication. The last term of (3) represents the loss of a copy of a file with two copies.

Relation (2) can be rewritten as

$$X_0^N(t) = X_0^N(0) + \mu \int_0^t X_1^N(u) du + M_0^N(t), \tag{4}$$

where $(M_0^N(t))$ is the martingale defined by

$$M_0^N(t) = \sum_{i=1}^{+\infty} \int_0^t \mathbf{1}_{\{i \leq X_1^N(u-)\}} [\mathcal{N}_{\mu,i}(du) - \mu du].$$

Its increasing process is given by

$$\langle M_0^N(t) \rangle = \mu \int_0^t X_1^N(u) du.$$

Since $X_1^N(u) \leq F_N$, there exists some constant C_0 such that

$$\mathbb{E}(M_0^N(t)^2) = \mathbb{E}(\langle M_0^N(t) \rangle) \leq C_0 N t$$

holds for all $t \geq 0$ and $N \geq 1$.

Similarly, if f is in $C_c(\mathbb{N})$, the set of functions with finite support on \mathbb{N} , (3) gives

$$\begin{aligned} f(X_1^N(t)) = & f(X_1^N(0)) + \mu \int_0^t [f(X_1^N(u) - 1) - f(X_1^N(u))] X_1^N(u) du \\ & + N \int_0^t \Omega \left[\frac{F_N}{N} - \frac{X_1^N(u) + X_0^N(u)}{N} \right] (f)(X_1^N(u)) du + M_1^N(t), \end{aligned} \tag{5}$$

where, for $y \geq 0$, $\Omega[y]$ is the functional operator defined by

$$\Omega[y](f)(x) = 2\mu y(f(x + 1) - f(x)) + \lambda \mathbf{1}_{\{x > 0\}}(f(x - 1) - f(x)), \quad x \in \mathbb{N},$$

and $(M_1^N(t))$ is a martingale such that, for some constant C_1 ,

$$\mathbb{E}(M_1^N(t)^2) \leq C_1 N \|f\|_{\infty, F_N} t$$

holds for all $t \geq 0$ and $N \geq 1$, where $\|f\|_{\infty, F_N} = \max\{|f(x)| : 0 \leq x \leq F_N\}$.

3. Fluid limits

In this section we prove that a significant fraction of files is lost quickly if the network is not correctly dimensioned. When the ratio $\rho = \lambda/\mu$ is less than 2β , for a large N , the fraction of files with two copies at time t , $(F_N - X_0^N(t) - X_1^N(t))/N$, is close to $\rho/2$ if t is large enough. Consequently, $(\beta - \rho/2)N$ files are lost and the network stabilizes for a subset of files with two copies whose cardinality is of the order $\rho/2$.

When $\rho = 2\beta$, this is the critical case which is analyzed in Section 4, where we prove that the number of lost files is of the order \sqrt{N} . When $\rho > 2\beta$, the proportion of lost files is 0 at the fluid level. This case is investigated in Section 5.

Theorem 1. (Fluid equations.) *If $\lim_{N \rightarrow +\infty} F_N/N = \beta$ then the sequence of processes $(X_0^N(t)/N, X_1^N(t)/N)$ converges in distribution to*

$$\begin{aligned} & [(\beta - \frac{1}{2}\rho)(1 - 2e^{-\mu t} + e^{-2\mu t}), (2\beta - \rho)(e^{-\mu t} - e^{-2\mu t})] \quad \text{if } \rho \leq 2\beta, \\ & (0, 0) \quad \text{if } \rho > 2\beta. \end{aligned}$$

Proof. Equations (4) and (5), with the function $f \equiv \text{Id}$ on $[0, F_N]$, can be written as

$$\frac{X_0^N(t)}{N} = \mu \int_0^t \frac{X_1^N(u)}{N} du + \frac{M_0^N(t)}{N}, \tag{6}$$

$$\begin{aligned} \frac{X_1^N(t)}{N} &= 2\mu \int_0^t \left(\frac{F_N}{N} - \frac{X_1^N(u) + X_0^N(u)}{N} \right) du - \lambda t - \mu \int_0^t \frac{X_1^N(u)}{N} du + \frac{M_1^N(t)}{N} \\ &+ \lambda \int_0^t \mathbf{1}_{\{X_1^N(u)=0\}} du. \end{aligned} \tag{7}$$

Doob’s inequality and the bounds on the second moments of the associated martingales show that, for $i = 0, 1$ and $t \geq 0$,

$$\mathbb{P} \left(\sup_{0 \leq s \leq t} \frac{M_i^N(s)}{N} \geq \varepsilon \right) \leq \frac{1}{\varepsilon^2} \mathbb{E}(M_i(t)^2) \leq \frac{1}{N} \frac{C_i t}{\varepsilon^2}.$$

Therefore, the two sequences of processes $(M_0^N(t)/N)$ and $(M_1^N(t)/N)$ converge in distribution to 0 uniformly on compact sets.

For $T > 0, \delta > 0$, and $i = 0, 1$, define $w_{X_i^N}^T(\delta)$ to be the modulus of continuity of the process $(X_i^N(t))$ on the interval $[0, T]$:

$$w_{X_i^N}^T(\delta) = \sup_{0 \leq s \leq t \leq T, |t-s| \leq \delta} |X_i^N(t) - X_i^N(s)|. \tag{8}$$

Using the fact that, for some constant $C, X_i^N(t) \leq F_N \leq CN$ for all $N \in \mathbb{N}$ and $t \geq 0$, it follows from the above equations and the convergence of the martingales to 0 that, for any $\varepsilon > 0$ and $\eta > 0$, there exists $\delta > 0$ such that the relation $\mathbb{P}(w_{X_i^N}^T(\delta) \geq \eta) \leq \varepsilon$ holds for all N .

The above implies that the sequence of stochastic processes $(X_0^N(t)/N, X_1^N(t)/N)$ is tight. See [3, Theorem 15.1] for example. We denote by $(x_0(t), x_1(t))$ a limiting value for some subsequence (N_k) . From (6), we obtain the relation

$$x_0(t) = \mu \int_0^t x_1(u) du.$$

Define

$$Z^N(t) = \mu \int_0^t \left(\frac{2F_N}{N} - \frac{3X_1^N(u)}{N} - \frac{2X_0^N(u)}{N} \right) du - \lambda t + \frac{M_1^N(t)}{N}. \tag{9}$$

Equation (7) can also be interpreted as the unique solution

$$(X_Z^N(t), R_Z^N(t)) := \left(\frac{X_1^N(t)}{N}, \lambda \int_0^t \mathbf{1}_{\{X_1^N(u)=0\}} du \right) \tag{10}$$

of the Skorokhod problem associated with the process $(Z^N(t))$. See Appendix A for a definition.

The sequence $(Z^{N_k}(t))$ converges in distribution and, by the continuous mapping theorem,

$$\begin{aligned} \lim_{k \rightarrow +\infty} (Z^{N_k}(t)) &= y(t) \\ &:= \mu \int_0^t (2\beta - 2x_0(u) - 3x_1(u)) du - \lambda t \\ &= (2\mu\beta - \lambda)t - \mu \int_0^t \left(3x_1(u) + 2\mu \int_0^u x_1(v) dv \right) du. \end{aligned} \tag{11}$$

As the solutions of Skorokhod problems are continuous with respect to the process $(Z^N(t))$ (see Appendix D of [31] for example), $(X_Z^N(t), R_Z^N(t))$ converges in distribution to the solution $(x_y(t), r_y(t))$ of the Skorokhod problem associated with $(y(t))$. Since $x_y(t) = x_1(t)$ and $y(t) = F(x_1(t))$ with

$$F(x)(t) = (2\mu\beta - \lambda)t - \mu \int_0^t \left(3x(u) + 2\mu \int_0^u x(v) dv \right) du,$$

the process $(x_1(t))$ is therefore the first coordinate of the solution of the generalized Skorokhod problem associated with the functional F . See Appendix A. Proposition 3 shows that such a solution exists and is unique. This implies that there is a unique, deterministic limiting value for the sequence $(X_0^N(t)/N, X_1^N(t)/N)$. It is easy to check that the explicit expressions for $(x_0(t))$ and $(x_1(t))$ given in the statement of the theorem are indeed the solutions of the generalized Skorokhod problem. The convergence in distribution is therefore established.

4. The critical case

To complete the picture of the overloaded network $\rho \leq 2\beta$, we consider the critical case $\rho = 2\beta$. As will be seen, the convergence result can be expressed in terms of a reflected stochastic differential equation. In Appendix A we present the corresponding definition, and a result of existence and uniqueness.

Theorem 2. *Let $\lambda/\mu = 2\beta$,*

$$\lim_{N \rightarrow +\infty} \frac{1}{\sqrt{N}} \left(F_N - N \frac{\rho}{2} \right) = \gamma \quad \text{for some } \gamma \in \mathbb{R}, \quad \lim_{N \rightarrow +\infty} \frac{X_1^N(0)}{\sqrt{N}} = y,$$

and $X_0^N(0) = 0$. Then the relation

$$\lim_{N \rightarrow +\infty} \left(\frac{X_0^N(t)}{\sqrt{N}}, \frac{X_1^N(t)}{\sqrt{N}} \right) = \left(\mu \int_0^t Y(u) du, Y(t) \right)$$

holds for the convergence in distribution, where $(Y(t))$ is the solution starting at y of the

stochastic differential equation

$$dY(t) = \sqrt{2\lambda} dB(t) + \mu \left(2\gamma - 3Y(t) - 2\mu \int_0^t Y(u) du \right) dt \tag{12}$$

reflected at 0, i.e. with the constraint that $Y(t) \geq 0$ for all $t \geq 0$. The process $(B(t))$ is a standard Brownian motion on \mathbb{R} .

The solution of stochastic differential equation (12) is non-Markovian due to the integral term in the drift.

Proof of Theorem 2. Equations (4) and (5), with the function $f \equiv \text{Id}$ on $[0, F_N]$, can be written as

$$\bar{X}_0^N(t) := \frac{X_0^N(t)}{\sqrt{N}} = \mu \int_0^t \frac{X_1^N(u)}{\sqrt{N}} du + \frac{M_0^N(t)}{\sqrt{N}}, \tag{13}$$

$$\begin{aligned} \bar{X}_1^N(t) &:= \frac{X_1^N(t)}{\sqrt{N}} \\ &= \frac{X_1^N(0)}{\sqrt{N}} + 2\mu \int_0^t \left(\gamma_N - \frac{X_1^N(u)}{\sqrt{N}} - \frac{X_0^N(u)}{\sqrt{N}} \right) du - \mu \int_0^t \frac{X_1^N(u)}{\sqrt{N}} du + \frac{M_1^N(t)}{\sqrt{N}} \\ &\quad + \lambda\sqrt{N} \int_0^t \mathbf{1}_{\{X_1^N(u)=0\}} du, \end{aligned} \tag{14}$$

with $\gamma_N = (F_N - N\rho/2)/\sqrt{N}$. With the same notation as used in Section 2, the martingales $(M_0^N(t))$ and $(M_1^N(t))$ are given by

$$\begin{aligned} M_0^N(t) &= \sum_{i=1}^{+\infty} \int_0^t \mathbf{1}_{\{i \leq X_1^N(u-)\}} [\mathcal{N}_{\mu,i}(du) - \mu du], \\ M_1^N(t) &= \sum_{i=1}^{+\infty} \int_0^t \mathbf{1}_{\{i \leq F_N - X_1^N(u-) - X_0^N(u-)\}} [\mathcal{N}_{2\mu}(du) - 2\mu du] - M_0^N(t) \\ &\quad - \int_0^t \mathbf{1}_{\{X_1^N(u) > 0\}} [\mathcal{N}_{\lambda N}(du) - \lambda N du]. \end{aligned}$$

Their increasing processes are given by

$$\begin{aligned} \left\langle \frac{1}{\sqrt{N}} M_0^N \right\rangle(t) &= \mu \int_0^t \frac{X_1^N(u)}{N} du, \\ \left\langle \frac{1}{\sqrt{N}} M_1^N \right\rangle(t) &= 2\mu \int_0^t \left(\frac{F_N}{N} - \frac{X_1^N(u)}{N} - \frac{X_0^N(0)}{N} \right) du + \left\langle \frac{1}{\sqrt{N}} M_0^N \right\rangle(t) \\ &\quad + \lambda \int_0^t \mathbf{1}_{\{X_1^N(u) > 0\}} du. \end{aligned}$$

The last term on the right-hand side of the above equation is $(R^N(t))$, defined as in (10). It is the second component of the solution to the Skorokhod problem associated with the process $(Z^N(t))$ of (9). The sequence of processes $(Z^N(t))$ converges to $(y(t))$ defined in (11). In this case $(y(t))$ is identically 0; the solution of the corresponding Skorokhod problem associated with $(y(t))$ is therefore $(0, 0)$. The continuity properties of the solutions of the Skorokhod

problem imply that the process $(R^N(t))$ converges to 0. Consequently, by Theorem 1 we obtain the convergence in distribution

$$\lim_{N \rightarrow +\infty} \left(\int_0^t \mathbf{1}_{\{X_1^N(u)=0\}} \, du \right) = 0,$$

and, therefore,

$$\lim_{N \rightarrow +\infty} \left(\left\langle \frac{1}{\sqrt{N}} M_0^N \right\rangle(t) \right) = 0 \quad \text{and} \quad \lim_{N \rightarrow +\infty} \left(\left\langle \frac{1}{\sqrt{N}} M_1^N \right\rangle(t) \right) = 2\lambda t.$$

We deduce that $(\bar{M}_1^N(t)) := (M_1^N(t)/\sqrt{N})$ converges to $(\sqrt{2\lambda}B(t))$, where $(B(t))$ is a standard Brownian motion, and that $(\bar{M}_0^N(t)) := (M_0^N/\sqrt{N})$ converges to 0. See [9, Theorem 1.4, p. 339] for example.

We now prove that the processes

$$(\bar{X}_0^N(t)) := \left(\frac{X_0^N(t)}{\sqrt{N}} \right) \quad \text{and} \quad (\bar{X}_1^N(t)) := \left(\frac{X_1^N(t)}{\sqrt{N}} \right)$$

are tight. If $(h(t))$ is a function on \mathbb{R}_+ , we define

$$\|h\|_{\infty,t} = \sup_{0 \leq s \leq t} |h(s)|$$

and let $w_h^t(\cdot)$ be the modulus of continuity of h defined in (8). Equation (13) gives, for $0 \leq t \leq T$,

$$\|\bar{X}_0^N\|_{\infty,t} \leq \|\bar{M}_0^N\|_{\infty,T} + \mu \int_0^t \|\bar{X}_1^N\|_{\infty,u} \, du.$$

Equation (14) shows that $(X_1^N(t)/\sqrt{N})$ is the first coordinate of the solution of the Skorokhod problem associated with $(Z_1^N(t))$ defined by

$$Z_1^N(t) := y_N + \mu \int_0^t \left(2\gamma_N - 3 \frac{X_1^N(u)}{\sqrt{N}} + 2 \frac{X_0^N(u)}{\sqrt{N}} \right) du + \frac{M_1^N(t)}{\sqrt{N}},$$

with $y_N = X_1^N(0)/\sqrt{N}$. Using the explicit representation of the solution of a Skorokhod problem in dimension 1, we have

$$\|\bar{X}_1^N\|_{\infty,t} \leq 2\|Z_1^N\|_{\infty,t} \quad \text{for } 0 \leq t \leq T;$$

see Appendix D of [31] for example. Then

$$\begin{aligned} \|\bar{X}_1^N\|_{\infty,t} &\leq 2y_N + 4\mu\gamma_N T + 2\|\bar{M}_1^N\|_{\infty,T} + 4\mu \int_0^t (\|\bar{X}_0^N\|_{\infty,u} + \|\bar{X}_1^N\|_{\infty,u}) \, du \\ &\leq U^N(T) + (4 + \mu T)\mu \int_0^t \|\bar{X}_1^N\|_{\infty,u} \, du, \end{aligned}$$

with $U^N(T) := 2y_N + 4\mu\gamma_N T + 2\|\bar{M}_1^N\|_{\infty,T} + 4\mu T\|\bar{M}_0^N\|_{\infty,T}$.

It follows from Gronwall’s inequality that the relation $\|\bar{X}_1^N\|_{\infty,t} \leq U^N(T) \exp((4 + \mu T)\mu t)$ holds for $0 \leq t \leq T$, and, consequently,

$$\|\bar{X}_0^N\|_{\infty,t} \leq \mu T U^N(T) e^{(4+\mu T)\mu t} + \|\bar{M}_0^N\|_{\infty,T}.$$

The convergence of martingales shows that the two sequences of random variables $(U^N(T))$ and $\|\bar{M}_0^N\|_{\infty,T}$ converge in distribution. Consequently, for $\varepsilon > 0$, there exists some $K > 0$

such that, for $i = 0, 1$ and all $N \geq 0$,

$$\mathbb{P}(\|\bar{X}_i^N\|_{\infty,t} > K) \leq \varepsilon.$$

If $\eta > 0$, there exist N_0 and sufficiently small δ such that, for all $N \geq N_0$,

$$2\mu\delta T(\gamma_N + 2K) < \frac{\eta}{2} \quad \text{and} \quad \mathbb{P}(w_{M_1^N}^T \geq \eta) \leq \varepsilon.$$

The last relation follows from the fact that the sequence $(\bar{M}_1^N(t))$ converges in distribution to a continuous process. We finally obtain

$$\begin{aligned} \mathbb{P}(w_{Z_1^N}^T(\delta) \geq \eta) &\leq \mathbb{P}(2\mu\delta T[\gamma_N + \|\bar{X}_0^N\|_{\infty,T} + \|\bar{X}_1^N\|_{\infty,T}] + w_{M_1^N}^T(\delta) \geq \eta) \\ &\leq \mathbb{P}(\|\bar{X}_0^N\|_{\infty,T} \geq K) + \mathbb{P}(\|\bar{X}_1^N\|_{\infty,T} \geq K) + \mathbb{P}\left(w_{M_1^N}^T(\delta) \geq \frac{\eta}{2}\right) \\ &\leq 3\varepsilon. \end{aligned}$$

The sequence $(Z_1^N(t))$ is therefore tight. By continuity of the solution of the Skorokhod problem, the same property holds for $(X_1^N(t)/\sqrt{N})$ and, consequently, for $(X_0^N(t)/\sqrt{N})$.

Let $(Y_0(t), Y_1(t))$ be a limit of the subsequence $[(X_0^{N_k}(t)/\sqrt{N_k}, X_1^{N_k}(t)/\sqrt{N_k})]$. By (13) and (14), we find that

$$\left(\frac{X_1^{N_k}(t)}{\sqrt{N_k}}, \lambda\sqrt{N} \int_0^t \mathbf{1}_{\{X_1^N(u)=0\}} du\right)$$

converges in distribution to the solution of the Skorokhod problem associated with the process

$$\left(y + \sqrt{2\lambda}B(t) + \mu \int_0^t \left(2\gamma - 3Y_1(u) - 2\mu \int_0^u Y_1(v) dv\right) du\right).$$

We conclude that $(Y_1(t))$ is the solution of the generalized Skorokhod problem for the functional F defined by

$$F(h)(t) = y + \sqrt{2\lambda}B(t) + \mu \int_0^t \left(2\gamma - 3h(u) - 2\mu \int_0^u h(v) dv\right) du.$$

Proposition 3 shows that there is a unique solution $(Y_1(t))$ and, consequently, a unique limit $(Y_0(t), Y_1(t))$. This completes the proof of the theorem.

5. The time scales of the stable network

The asymptotic properties of the network are investigated under the condition $\rho = \lambda/\mu > 2\beta$. In Section 3 we showed that in this case the system is stable at the fluid level, i.e. that the fraction of lost files is 0. Of course, this does not change the fact that the system is still transient with the absorbing state $(F_N, 0)$. To understand how the system reaches this state, we must consider the following three interesting time scales.

- The slow time scale: $t \rightarrow t/N$.
- The normal time scale: $t \rightarrow t$.
- The linear time scale: $t \rightarrow Nt$.

We successively investigate each of these time scales in this section. The following elementary lemma will be used throughout the section.

Lemma 1. *If $\rho = \lambda/\mu > 2\beta$ for any $\beta_0 > \beta$ such that $\lambda/\mu > 2\beta_0$, $\varepsilon > 0$, $\eta > 0$, and $T > 0$, there exists $N_0 \in \mathbb{N}$ such that*

(i) *(coupling) there exists a probability space where the relation*

$$X_1^N(t) \leq L_{\beta_0}(Nt) \quad \text{for all } t \geq 0$$

holds for all $N \geq N_0$ and $t \geq 0$, where $(L_{\beta_0}(t))$ is the process of the number of customers in an ergodic M/M/1 queue with arrival rate $2\mu\beta_0$ and service rate λ , with initial condition $L_{\beta_0}(0) = X_1^N(0)$,

(ii) *the relation*

$$\mathbb{P} \left[\sup_{\substack{0 \leq s, t \leq T \\ |t-s| \leq \delta}} \frac{1}{N} \int_{sN}^{tN} L_{\beta_0}(u) \, du > \eta \right] \leq \varepsilon$$

holds.

Proof. There exist some $\beta_0 \geq \beta$ and $N_0 \geq 1$ such that $\lambda > 2\beta_0\mu$ and $F_N \leq N\beta_0$ for $N \geq N_0$. It is enough to take the M/M/1 queue with arrival rate $2\mu\beta_0$ and service rate λ .

Denote by \mathcal{A} the event on the left-hand side of the relation in (ii). If, for $x \in \mathbb{N}$, τ_x denotes the hitting time of x by the process $(L_{\beta_0}(t))$, for $\delta < \frac{1}{2}$, we have

$$\mathbb{P}(\mathcal{A}) \leq \mathbb{P}(\tau_{\lfloor \eta N \rfloor} \leq NT).$$

By ergodicity of this process and Proposition 5.11 of [31] for example, there exist some $0 < \alpha < 1$ such that the sequence $(\alpha^N \tau_{\lfloor N\eta \rfloor})$ converges in distribution. The right-hand side of the above relation is thus arbitrarily small as N becomes large.

5.1. The slow time scale

A description of the asymptotic behavior for the slow time scale is briefly presented. The following proposition shows in particular that, on the slow time scale, with probability 1, no file is lost in the limit.

Proposition 1. *The sequence of processes $(X_0^N(t/N), X_1^N(t/N))$ converges in distribution to the process $(0, L_\beta(t))$, where $(L_\beta(t))$ is the process of the number of jobs in an M/M/1 queue with arrival rate $2\mu\beta$ and service rate λ .*

Proof. From (1), we can see that the Q -matrix of the process on the slow time scale $(X_0^N(t/N), X_1^N(t/N))$ has the following asymptotic expansion:

$$\begin{aligned} q^N(x, x + e_1) &= 2\mu \frac{1}{N} (F_N - x_0 - x_1) \rightarrow 2\mu\beta, \\ q^N(x, x - e_1) &= \lambda \mathbf{1}_{\{x_1 > 0\}}, \\ q^N(x, x - e_1 + e_0) &= \mu \frac{x_1}{N} \rightarrow 0. \end{aligned}$$

For $T > 0$, the above lemma shows that, for sufficiently large K , $(X_1^N(t/N))$ is bounded by K on the time interval $[0, T]$ with high probability and for any $N \geq 0$. In particular, for sufficiently large M , the number of jumps of the process $(X_1^N(t/N), 0 \leq t \leq T)$ is less than

M with high probability for all N . The convergence of the successive instants of jumps is then a direct consequence of the convergence of the Q -matrix. The convergence of the processes on the time interval $[0, T]$ follows from the definition of the metrics defining the Skorokhod topology on right-continuous functions with left limits. See [9, Proposition 5.2] for example.

5.2. The normal time scale

We show that, on the normal time scale, the stability does not only hold on the fluid level: almost surely, there is a finite number of losses in any finite time interval; more precisely, losses occur as a Poisson process (see Theorem 3). The capacity λN of the network is thus able to maintain an almost complete set of files. The following proposition shows in particular that the number of definitive losses at time $t > 0$ is finite with a Poisson distribution.

Theorem 3. *If $\rho = \lambda/\mu > 2\beta$,*

- (i) *the sequence of processes $(X_0^N(t))$ converges in distribution to a Poisson point process on \mathbb{R}_+ with rate $2\mu\beta/(\rho - 2\beta)$,*
- (ii) *for $t > 0$, as N goes to ∞ , the random variable $X_1^N(t)$ converges in distribution to a geometric distribution with parameter $2\beta/\rho$.*

The second convergence is for the marginal distribution of $(X_1^N(s))$ at time t . One cannot expect a convergence in distribution of the sequence of processes $(X_1^N(t))$. Indeed, since the sequence of processes $(X_1^N(t/N))$ converges in distribution to the law of the M/M/1 process $(L_\beta(t))$ for $0 \leq s < t$, the distributions of $(X_1^N(s), X_1^N(t))$ and $(L_\beta(Ns), L_\beta(Nt))$ are close. Between times Ns and Nt , the M/M/1 ‘forgets’ its location at time Ns (just because it hits 0 with probability 1), so when N goes to ∞ , the pair $(X_1^N(s), X_1^N(t))$ converges in distribution to the distribution of two independent geometric distributions.

Proof of Theorem 3. Define

$$\eta_N(t) := \int_0^t X_1^N(u) \, du.$$

For $0 \leq s \leq t$, Lemma 1 gives

$$\eta_N(t) - \eta_N(s) = \int_s^t X_1^N(u) \, du \leq \int_s^t L_{\beta_0}(Nu) \, du = \frac{1}{N} \int_{Ns}^{Nt} L_{\beta_0}(u) \, du.$$

It follows from the criteria of the modulus of continuity and Lemma 1 that the sequence of processes $(\eta^N(t))$ is tight. The above inequality and the ergodic theorem applied to the ergodic Markov process $(L_{\beta_0}(t))$ also show that, almost surely,

$$\limsup_{N \rightarrow +\infty} \int_0^t X_1^N(u) \, du \leq \frac{2\beta_0}{\rho - 2\beta_0} t. \tag{15}$$

For fixed $T > 0$ and $K > 0$,

$$\begin{aligned} \mathbb{P}(X_0^N(T) \geq K) &\leq \mathbb{P}\left(\mu \int_0^T X_1^N(u) \, du \geq \frac{K}{2}\right) + \mathbb{P}\left(M_0^N(T) \geq \frac{K}{2}\right) \\ &\leq \mathbb{P}\left(\mu \int_0^T X_1^N(u) \, du \geq \frac{K}{2}\right) + \frac{4}{K^2} \mathbb{E}\left(\mu \int_0^T X_1^N(u) \, du\right). \end{aligned}$$

We can thus choose K so that $\mathbb{P}(X_0^N(T) \geq K) \leq \varepsilon$ holds for $N \geq N_0$ for some $N_0 \in \mathbb{N}$. As in the proof of Lemma 1, for $\delta > 0$, there exist some $N_1 \in \mathbb{N}$ such that if $N \geq N_1$ then

$$\mathbb{P}\left(\sup_{0 \leq s \leq T} L_{\beta_0}(s) \geq \delta N\right) \leq \varepsilon.$$

In the same way as in the proof of Lemma 1, we can construct an M/M/1 process $(Z^N(t))$ whose arrival and service rates are respectively

$$2\mu\left(\frac{F_N}{N} - \frac{K}{N} - \delta\right) \quad \text{and} \quad \lambda,$$

and such that, on the event

$$\mathcal{A}_T := \left\{X_0^N(T) \leq K, \sup_{0 \leq t \leq NT} L_{\beta_0}(t) \leq \delta N\right\},$$

the relation $X_1^N(t) \geq Z^N(Nt)$ holds for all $t \leq T$. Hence, almost surely,

$$\liminf_{N \rightarrow +\infty} \eta^N(t) \geq \liminf_{N \rightarrow +\infty} \frac{1}{N} \int_0^{Nt} Z^N(u) \, du = \frac{2(\beta - \delta)}{\rho - 2(\beta - \delta)} t \tag{16}$$

holds on \mathcal{A}_T . By setting β_0 equal to β in (15) and δ equal to 0 in (16), it follows that the variable $\eta^N(t)$ converges almost surely to αt with $\alpha = 2\beta\mu/(\rho - 2\beta)$. Consequently, the tightness of the sequence of processes $(\eta^N(t))$ implies that it converges in distribution to (αt) .

Note that $t \mapsto X_0^N(t)$ can also be seen as a point process with jumps of size 1. By (4),

$$\left(X_0^N(t) - \mu \int_0^t X_1^N(u) \, du\right)$$

is a martingale with respect to the natural filtration of the associated Poisson processes. The random measure

$$\Lambda^N([0, t]) = \mu \int_0^t X_1^N(u) \, du$$

is a *compensator* of the point process $t \mapsto X_0^N(t)$; see [19]. It has therefore been shown that the sequence of compensators converges to the deterministic measure $\alpha \, dx$. Theorem 5.1 of [19], see also [4], gives the convergence in distribution of $(X_0^N(t))$ to a Poisson process with rate α .

In a similar way as in the proof of Proposition 1, through the convergence of the Q -matrix, the asymptotic distribution of $X_1^N(t)$ can easily be obtained by conditioning on the event $\{X_0^N(t) \leq K\}$ for large K and by using arbitrarily close M/M/1 processes at equilibrium as upper and lower stochastic bounds for $X_1^N(t)$. We omit the details.

5.3. The linear time scale $t \rightarrow Nt$

On the linear time scale, we will show that a fraction $\Psi(t)$ of the files is lost at time t . That is, the linear time scale gives a picture of the decay of the network. In the following, for a nonnegative measure Λ on some space H and a nonnegative measurable function f on H , the classical notation $\langle \Lambda, f \rangle$ will denote the integration of f with respect to the measure Λ .

For $N \geq 1$, the random measure μ_N on $\mathbb{N} \times \mathbb{R}_+$ is defined as, for a measurable function $g: \mathbb{N} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$,

$$\langle \mu_N, g \rangle = \int_{\mathbb{R}_+} g(X_1^N(Nt), t) \, dt.$$

Note that if $g(x, t) = h(x) \mathbf{1}_{[0, T]}(t)$ for $T > 0$ then

$$\langle \mu_N, g \rangle = \sum_{x \in \mathbb{N}} h(x) \frac{1}{N} \int_0^{NT} \mathbf{1}_{\{X_1^N(t)=x\}} dt.$$

Consequently, (μ_N) is a relatively compact sequence of random Radon measures on $\mathbb{N} \times \mathbb{R}_+$; see [7, Theorem 3.7.1, p. 56] for example. Note that the measure identically null can be a possible limit of this sequence.

In the following we fix (N_k) such that (μ_{N_k}) is a converging subsequence whose limit is ν . By taking a convenient probability space, we can assume that the convergence of (μ_{N_k}) holds almost surely for the weak convergence of Radon measures.

Since, for $N \geq 1$, μ_N is absolutely continuous with respect to the product of the counting measure on \mathbb{N} and Lebesgue measure on \mathbb{R}_+ , the same property holds for the limiting measure ν . Let $(x, t) \rightarrow \pi_t(x)$ denote its (random) density. We note that a version of $\pi_t(x)$ can be chosen such that the map $(\omega, x, t) \rightarrow \pi_t(x)(\omega)$ on the product of the probability space and $\mathbb{N} \times \mathbb{R}_+$ is measurable by taking $\pi_t(x)$ as a limit of measurable maps:

$$\pi_t(x) = \limsup_{s \rightarrow 0} \frac{1}{s} \nu(\{x\} \times [t, t + s]).$$

See Chapter 8 of [33] for example.

Proposition 2. *For the convergence in distribution of continuous processes,*

$$\lim_{k \rightarrow +\infty} \left(\frac{\mu}{N_k} \int_0^{N_k t} X_1^{N_k}(u) du \right) = (\Psi(t)) := \left(\mu \int_0^t \langle \pi_u, I \rangle du \right),$$

where $I(x) = x$ for $x \in \mathbb{N}$. Moreover, almost surely, for all $t \geq 0$,

$$\int_0^t \pi_u(\mathbb{N}) du = t.$$

The latter relation is crucial, as it shows that the masses of the measures μ_{N_k} for $k \geq 1$ do not vanish at ∞ . This property is sometimes absent from the proofs of SAPs, but is nevertheless essential to identify π_u as an invariant distribution of a Markov process.

Proof of Proposition 2. The criteria of the modulus of continuity is used to prove the tightness of

$$(\Psi_N(t)) := \left(\frac{\mu}{N} \int_0^{Nt} X_1^N(u) du \right).$$

By Lemma 1,

$$\Psi_N(t) - \Psi_N(s) = \frac{\mu}{N} \int_{Ns}^{Nt} X_1^N(u) du \leq \frac{\mu}{N^2} \int_{N^2s}^{N^2t} L_{\beta_0}(u) du.$$

As in the proof of Theorem 3, we conclude that the sequence of processes $(\Psi_N(t))$ is tight.

For $K > 0$ and $t \geq 0$, the almost-sure convergence of the measures (μ_{N_k}) gives the convergence

$$\lim_{k \rightarrow +\infty} \frac{1}{N_k} \int_0^{N_k t} X_1^{N_k}(u) \mathbf{1}_{[0, K]}(X_1^{N_k}(u)) du = \int_0^t \langle \pi_u, I \mathbf{1}_{[0, K]} \rangle du,$$

where $I(x) = x$. By again using Lemma 1 we obtain

$$\frac{1}{N_k} \int_0^{N_k t} X_1^{N_k}(u) \mathbf{1}_{\{X_1^{N_k}(u) \geq K\}} du \leq \frac{1}{N_k^2} \int_0^{N_k^2 t} L_{\beta_0}(u) \mathbf{1}_{\{L_{\beta_0}(u) \geq K\}} du,$$

and the ergodic theorem applied to $(L_{\beta_0}(t))$ shows that the last quantity converges in distribution to

$$(t\mathbb{E}(L_{\beta_0}(\infty) \mathbf{1}_{\{L_{\beta_0}(\infty) \geq K\}})),$$

where $L_{\beta_0}(\infty)$ is the limit in distribution of $(L_{\beta_0}(t))$, a geometrically distributed random variable. For $\varepsilon > 0$, K is chosen sufficiently large so that the last quantity is less than $\varepsilon/2$; consequently, if k is large enough, we have

$$\frac{1}{N_k} \int_0^{N_k t} X_1^{N_k}(u) \mathbf{1}_{\{X_1^{N_k} \geq K\}} du \leq \varepsilon.$$

We deduce that $(\Psi(t))$ is the only possible limiting process for $(\Psi_{N_k}(t))$. This proves the first part of the proposition.

For $K \geq 1$, the convergence of (μ_{N_k}) gives the relation

$$\lim_{k \rightarrow +\infty} \frac{1}{N_k} \int_0^{N_k t} \mathbf{1}_{\{X_1^{N_k}(u) \leq K\}} du = \nu([0, K] \times [0, t]) = \int_0^t \pi_u([0, K]) du.$$

By again using the stochastic domination for an ergodic M/M/1 queue,

$$\frac{1}{N_k} \int_0^{N_k t} \mathbf{1}_{\{L_{\beta_0}(u) \leq K\}} du \leq \frac{1}{N_k} \int_0^{N_k t} \mathbf{1}_{\{X_1^{N_k}(u) \leq K\}} du.$$

Letting k go to ∞ , we obtain, almost surely,

$$t\mathbb{P}(L_{\beta_0}(\infty) \leq K) \leq \int_0^t \pi_u([0, K]) du \leq \int_0^t \pi_u(\mathbb{N}) du.$$

If K goes to ∞ , the relation

$$\int_0^t \pi_s(\mathbb{N}) ds = t$$

holds for all $t \in \mathbb{N}$ and, consequently, for all $t \geq 0$. This completes the proof.

Theorem 4. (Rate of decay of the network.) *If $\rho = \lambda/\mu > 2\beta$ then, as N goes to ∞ , the process $(X_0^N(Nt)/N)$ converges to $(\Psi(t))$, where $\Psi(t)$ is the unique solution $y \in [0, \beta]$ of the equation*

$$\left(1 - \frac{y}{\beta}\right)^{\rho/2} e^{y+\mu t} = 1.$$

For $t \geq 0$, the process $(X_1^N(Nt + u), u > 0)$ converges in distribution to the stationary process of the number of jobs in an M/M/1 queue with service rate λ and arrival rate $2\mu(\beta - \Psi(t))$.

It is easily seen that the asymptotic expansion $\Psi(t) - \beta \sim \beta \exp(-2(\beta + \mu t)/\rho)$ holds as t goes to ∞ . The last part of the theorem states that, ‘around’ time Nt , the process X_1^N has a local equilibrium.

Proof of Theorem 4. Equation (5) gives, for $f \in C_c(\mathbb{N})$,

$$f(X_1^N(Nt)) - f(X_1^N(0)) - M_{f,1}^N(Nt) = N^2 \int_0^t \Omega[Y_N(u)](f)(X_1^N(Nu)) du + \mu N \int_0^t \Delta^-(f)(X_1^N(Nu))X_1^N(Nu) du, \tag{17}$$

where, using (4),

$$Y_N(u) = \frac{F_N}{N} - \frac{X_1^N(Nu)}{N} - \frac{M_0^N(Nu)}{N} - \frac{\mu}{N} \int_0^{Nu} X_1^N(v) dv,$$

and $\Delta^-(f)(x) = (f(x - 1) - f(x)) \mathbf{1}_{\{x \geq 1\}}$. The bound on the increasing process of the martingale $(M_{f,1}^N(t))$ at the end of Section 2, Doob’s inequality, and Lemma 1 show that the sequence of processes

$$\left(\frac{1}{N^2} \left[f(X_1^N(Nt)) - f(X_1^N(0)) - M_{f,1}^N(Nt) - \mu N \int_0^t (\Delta^-(f)(X_1^N(Nu))X_1^N(Nu) du) \right] \right)$$

converges to 0 for the topology of the uniform norm on compact sets.

By Lemma 1, we obtain

$$\frac{X_1^N(Nu)}{N} \leq \frac{L_{\beta_0}(N^2u)}{N};$$

hence, the sequence of processes $(X_1^N(Nu)/N)$ converges in distribution to 0.

The bound on the increasing process and Proposition 2 show that the sequence of processes $(Y_{N_k}(t))$ converges in distribution to $(\beta - \Psi(t))$. We deduce from (17) that the sequence of processes

$$\left(\int_0^t \Omega[\beta - \Psi(u)](f)(X_1^{N_k}(N_ku)) du \right) = \left(\int_{\mathbb{N} \times [0,t]} \Omega[\beta - \Psi(u)](f)(x) \mu_{N_k}(dx, du) \right)$$

converges to 0.

The convergence of the (μ_{N_k}) and Proposition 2 therefore show that, almost surely, the relations

$$\int_0^t \langle \pi_u, \Omega[\beta - \Psi(u)](f) \rangle du = 0 \quad \text{and} \quad \int_0^t \pi_u(\mathbb{N}) du = t$$

hold for all $t \geq 0$ and all functions $f \in C_c(\mathbb{N})$. Note that we have used the fact that $C_c(\mathbb{N})$ has a countable dense subset for the uniform norm.

If Δ is the subset of all real numbers $u \geq 0$ such that either of the relations

$$\pi_u(\mathbb{N}) \neq 1, \quad \langle \pi_u, \Omega[\beta - \Psi(u)](f) \rangle \neq 0 \quad \text{for some } f \in C_c(\mathbb{N}) \tag{18}$$

holds, then the Lebesgue measure of Δ is 0. Hence, if $u \notin \Delta$ then $\pi_u(\mathbb{N}) = 1$ and $\langle \pi_u, \Omega[\beta - \Psi(u)](f) \rangle = 0$ for all $f \in C_c(\mathbb{N})$. Since $\Omega[\beta - \Psi(u)]$ is the infinitesimal generator of an M/M/1 queue with arrival rate $2\mu(\beta - \Psi(u))$ and service rate λ , π_u is a geometric distribution on \mathbb{N} with parameter $2\mu(\beta - \Psi(u))/\lambda$.

From Proposition 2, for $t \geq 0$,

$$\Psi(t) = \mu \int_{[0,t] \setminus \Delta} \langle \pi_u, I \rangle du = \mu \int_0^t \frac{2\mu(\beta - \Psi(u))}{\lambda - 2\mu(\beta - \Psi(u))} du.$$

Straightforward calculus gives the relation

$$(\beta - \psi(u))^{\rho/2} e^{\psi(u)} = \beta^{\rho/2} e^{-\mu u}.$$

Since $2\beta < \rho$, it is easy to check that there is a unique $\Psi(u) < \beta$ satisfying the above equation. This completes the proof.

Corollary 1 below on the asymptotic behavior of $T_N(\delta)$, the first time when a fraction δ of the files has been lost, follows directly from Theorem 4.

Corollary 1. *If $\rho = \lambda/\mu > 2\beta$ then, for $N \geq 1$ and $\delta \in (0, 1)$,*

$$T_N(\delta) = \inf\{t \geq 0: X_0^N(t) \geq \delta F_N\}.$$

For the convergence in distribution,

$$\lim_{N \rightarrow +\infty} \frac{T_N(\delta)}{N} = \frac{1}{\mu} \left(-\frac{\rho}{2} \log(1 - \delta) - \delta\beta \right).$$

5.4. SAP

The quantity used to establish the theorem concerning the rate of decay of the network is

$$\frac{1}{N} \int_0^{Nt} X_1^N(u) \, du = \int_0^t X_1^N(Nu) \, du,$$

which is the average of lost files on $[0, Nt]$ scaled by N . If at ‘time’ Nu the number of remaining files (i.e. those not lost) is $N(\beta - \psi(u))$, the process $(X_1^N(Nu + v), v \geq 0)$ is an ergodic M/M/1 queue with input rate $2\mu(\beta - \psi(u))$ and service rate λ . The π_u of Proposition 2 is therefore the geometric distribution with parameter $2\mu(\beta - \psi(u))/\lambda$. The local equilibrium of (X_1^N) determines the evolution of the process X_0^N on the time scale $t \rightarrow Nt$.

This SAP phenomenon is already well known in the framework of deterministic dynamical systems; see [15]. In a stochastic context, an additional difficulty, sometimes neglected or underestimated in the literature, is of controlling the regularity properties of a family of invariant distributions indexed by the state of the slow process. This can be done through a kind of uniform control of some ergodic averages (see [14]) or by using a martingale representation of the associated Markov processes and a convenient Lyapunov function (see [24]). See also [22]. In any case, there are several delicate technical issues to address: a convenient tightness result for a set of random measures and the rate of convergence of ergodic averages. In our case this is obtained in Proposition 2 as a consequence of a simple monotonicity property.

Appendix A. Generalized Skorokhod problems

For the sake of self-containedness, in this section we briefly present the more or less classical material necessary to state and prove the convergence results used in this paper. The general theme concerns the rigorous definition of the solution of a stochastic differential equation constrained to stay in some domain, and also the proof of the existence and uniqueness of such a solution. See [1], [8], and [34], and [16] and [35] for the multidimensional context, and [29] for the more general context. See Appendix D of [31] for a brief account.

We first recall the classical definition of the Skorokhod problem. If $(Z(t))$ is some function of the set $\mathcal{D}(\mathbb{R}_+, \mathbb{R})$ of càdlàg functions defined on \mathbb{R}_+ , the pair of functions $[(X(t)), (R(t))]$ is said to be a solution of the Skorokhod problem associated with $(Z(t))$ whenever

1. $X(t) = Z(t) + R(t)$ for all $t \geq 0$,
2. $X(t) \geq 0$ for all $t \geq 0$,

3. $t \rightarrow R(t)$ is nondecreasing, $R(0) = 0$, and

$$\int_{\mathbb{R}_+} X(t) dR(t) = 0.$$

The generalization used in this paper corresponds to the case when $(Z(t))$ is itself a functional of $(X(t))$.

Definition 1. (*Generalized Skorokhod problem.*) If $G : \mathcal{D}(\mathbb{R}_+, \mathbb{R}) \rightarrow \mathcal{D}(\mathbb{R}_+, \mathbb{R})$ is a Borelian function, $((X(t)), (R(t)))$ is a solution of the generalized Skorokhod problem (GSP) associated with G if $((X(t)), (R(t)))$ is the solution of the Skorokhod problem associated with $G(X)$; in particular, for all $t \geq 0$,

$$X(t) = G(X)(t) + R(t) \quad \text{and} \quad \int_{\mathbb{R}_+} X(t) dR(t) = 0.$$

The classical Skorokhod problem described above corresponds to the case when the functional G is constant and equal to $(Z(t))$. If we take

$$G(x)(t) = \int_0^t \sigma(x(u)) dB(u) + \int_0^t \delta(x(u)) du,$$

where $(B(t))$ is a standard Brownian motion and σ and δ are Lipschitz functions on \mathbb{R} , the first coordinate $(X(t))$ of a possible solution to the corresponding GSP can be described as the solution of the stochastic differential equation

$$dX(t) = \sigma(X(t)) dB(t) + \delta(X(t)) dt$$

reflected at 0.

Proposition 3. *If $G : \mathcal{D}(\mathbb{R}_+, \mathbb{R}) \rightarrow \mathcal{D}(\mathbb{R}_+, \mathbb{R})$ is such that, for any $T > 0$, there exists a constant C_T such that, for all $(x(t)) \in \mathcal{D}(\mathbb{R}_+, \mathbb{R})$ and $0 \leq t \leq T$,*

$$\sup_{0 \leq s \leq t} |G(x)(s) - G(y)(s)| \leq C_T \int_0^t |x(u) - y(u)| du, \tag{19}$$

then there exists a unique solution to the GSP associated with the functional G and the matrix P .

Proof. Define the sequence $(X_N(t))$ by induction: $(X^0(t), R^0(t)) = 0$ and, for $N \geq 1$, (X^{N+1}, R^{N+1}) is the solution of the Skorokhod problem (SP) associated with $G(X^N)$, in particular,

$$X^{N+1}(t) = G(X^N)(t) + R^{N+1}(t) \quad \text{and} \quad \int_{\mathbb{R}_+} X^{N+1}(u) dR^{N+1}(u) = 0.$$

The existence of such a solution is guaranteed as well as the Lipschitz property of the solutions of a classical SP (see Proposition D.4 of [31]); this gives the existence of some constant K_T such that, for all $N \geq 1$ and $0 \leq t \leq T$,

$$\|X^{N+1} - X^N\|_{\infty,t} \leq K_T \|G(X^N) - G(X^{N-1})\|_{\infty,t},$$

where $\|h\|_{\infty,T} = \sup\{|h(s)| : 0 \leq s \leq T\}$. From (19), this implies that

$$\|X^{N+1} - X^N\|_{\infty,t} \leq \alpha \int_0^t \|X^N - X^{N-1}\|_{\infty,u} du,$$

with $\alpha = K_T C_T$. The iteration of the last relation yields the inequality

$$\|X^{N+1} - X^N\|_{\infty, t} \leq \frac{\alpha^N t^{N-1}}{(N-1)!} \int_0^t \|X^1\|_{\infty, u} \, du, \quad 0 \leq t \leq T.$$

We conclude that the sequence $(X^N(t))$ converges uniformly on compact sets and, consequently, so does the sequence $(R^N(t))$. Let $(X(t))$ and $(R(t))$ be the limit of these sequences. By continuity of the SP, the pair $((X(t)), (R(t)))$ is the solution of the SP associated with $G(X)$, and, hence, a solution of the GSP associated with G .

Uniqueness. Let $(Y(t))$ be another solution of the GSP associated with G . In the same way as before, we obtain, by induction, for $0 \leq t \leq T$,

$$\|X - Y\|_{\infty, t} \leq \frac{(\alpha t)^N}{N!} \int_0^t \|X - Y\|_{\infty, u} \, du.$$

Letting N go to ∞ , we conclude that $X = Y$. This completes the proof.

References

- [1] ANDERSON, R. F. AND OREY, S. (1976). Small random perturbation of dynamical systems with reflecting boundary. *Nagoya Math. J.* **60**, 189–216.
- [2] ARTALEJO, J. R. AND GÓMEZ-CORRAL, A. (2008). *Retrial Queueing Systems. A Computational Approach*. Springer, Berlin.
- [3] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*, 2nd edn. John Wiley, New York.
- [4] BROWN, T. (1978). A martingale approach to the Poisson convergence of simple point processes. *Ann. Prob.* **6**, 615–628.
- [5] CHUN, B.-G. *et al.* (2006). Efficient replica maintenance for distributed storage systems. In *Proc. NSDI*, IEEE, pp. 45–58.
- [6] DARROCH, J. N. AND SENETA E. (1965). On quasi-stationary distributions in absorbing discrete-time finite Markov chains. *J. Appl. Prob.* **2**, 88–100.
- [7] DAWSON, D. A. (1993). Measure-Valued Markov Processes. In *École d'Été de Probabilités de Saint-Flour XXI—1991* (Lecture Notes Math. **1541**), Springer, Berlin, pp. 1–260.
- [8] EL KAROUI, N. AND CHALEYAT-MAUREL, M. (1978). *Temps Locaux*, vol. 52–53, ch. Un problème de réflexion et ses applications au temps local et aux équations différentielles stochastiques sur \mathbb{R} -cas continu. In *Astérisque* (Temps Locaux 52–53), Société Mathématique de France, pp. 117–144.
- [9] ETHIER, S. N. AND KURTZ, T. G. (1986). *Markov Processes: Characterization and Convergence*. John Wiley, New York.
- [10] FALIN, G. I. AND TEMPLETON, J. G. C. (1997). *Retrials Queues*. Chapman & Hall, London.
- [11] FERRARI, P. A., KESTEN, H., MARTINEZ, S. AND PICCO, P. (1995). Existence of quasi-stationary distributions. A renewal dynamical approach. *Ann. Prob.* **23**, 501–521.
- [12] FEUILLET, M. (2012). On the flow-level stability of data networks without congestion control: the case of linear networks and upstream trees. *Queueing Systems* **70**, 105–143.
- [13] FEUILLET, M. AND ROBERT, P. (2012). On the transient behavior of Ehrenfest and Engset processes. *Adv. Appl. Prob.* **44**, 562–582.
- [14] FREIDLIN, M. I. AND WENTZELL, A. D. (1998). *Random Perturbations of Dynamical Systems*, 2nd edn. Springer, New York.
- [15] GUCKENHEIMER, J. AND HOLMES, P. (1990). Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. (*Appl. Math. Sci.* **42**). Springer, New York.
- [16] HARRISON, J. M. AND REIMAN, M. I. (1981). Reflected Brownian motion on an orthant. *Ann. Prob.* **9**, 302–308.
- [17] HAS' MINSKIĬ, R. Z. (1980). *Stochastic Stability of Differential Equations*. Sijthoff & Noordhoff, Alphen aan den Rijn, Germantown, MD.
- [18] HUNT, P. J. AND KURTZ, T. G (1994). Large loss networks. *Stoch. Process. Appl.* **53**, 363–378.
- [19] KASAHARA, Y. AND WATANABE, S. (1986). Limit theorems for point processes and their functionals. *J. Math. Soc. Japan* **38**, 543–574.
- [20] KELLY, F. P. (1991). Loss networks. *Ann. Appl. Prob.* **1**, 319–378.
- [21] KING, P. J. B. (1990). *Computer and Communication Systems Performance Modelling*. Prentice Hall, London.

- [22] KURTZ, T. G. (1992). Averaging for Martingale Problems and Stochastic Approximation. In *Applied Stochastic Analysis* (Lecture Notes Control Inf. Sci. **177**), Springer, Berlin, pp. 186–209.
- [23] LEGTCHENKO, S., MONNET, S., SENS, P. AND MULLER, G. (2011). RelaxDHT: a churn-resilient replication strategy for peer-to-peer distributed hash-tables. *ACM Trans. Autonomous Adaptive Systems* **7**, 28.
- [24] PAPANICOLAOU, G. C., STROOCK, D. AND VARADHAN, S. R. S. (1977). Martingale approach to some limit theorems. In *Papers from the Duke Turbulence Conference* (Duke University, Durham, NC, 1976), Paper number 6.
- [25] PAVLIOTIS, G. A. AND STUART, A. M. (2008). *Multiscale Methods. Averaging and Homogenization*. (Texts Appl. Math. **53**). Springer, New York.
- [26] PERRY, O. AND WHITT, W. (2011). An ODE for an overloaded X model involving a stochastic averaging principle. *Stoch. Systems* **1**, 59–108.
- [27] PICCONI, F., BAYNAT, B. AND SENS, P. (2007). An Analytical Estimation of Durability in DHTs. In *Distributed Computing and Internet Technology* (Lecture Notes Comput. Sci. **4882**), Springer, Berlin, pp. 184–196.
- [28] RAMABHADHRAN, S. AND PASQUALE, J. (2006). Analysis of long-running replicated systems. In *Proc. INFOCOM 2006*, IEEE, pp. 1–9.
- [29] RAMANAN, K. (2006). Reflected diffusions defined via the extended Skorokhod map. *Electron. J. Probab.* **11**, 934–992.
- [30] RHEA, S. *et al.* (2005). OpenDHT: a public DHT service and its uses. In *Proc. SIGCOMM '05*, ACM, New York, pp. 73–84.
- [31] ROBERT, P. (2003). *Stochastic Networks and Queues*. (Appl. Math. (New York) **52**). Springer, Berlin.
- [32] ROWSTRON, A. AND DRUSCHEL, P. (2001). Storage management and caching in PAST, a large-scale, persistent peer-to-peer storage utility. In *Proc. SOSP'01*, ACM, New York, pp. 188–201.
- [33] RUDIN, W. (1987). *Real and Complex Analysis*, 3rd edn. McGraw-Hill, New York.
- [34] SKOROKHOD, A. V. (1962). Stochastic equations for diffusion processes in a bounded region. II. *Theory Probab. Appl.* **7**, 3–23.
- [35] TAYLOR, L. M. AND WILLIAMS, R. J. (1993). Existence and uniqueness of semimartingale reflecting Brownian motions in an orthant. *Prob. Theory Relat. Fields* **96**, 283–317.