STABILITY PROPERTIES OF
LINEAR FILE-SHARING NETWORKS

Lasse Leskelä　　　　Philippe Robert　　　　Florian Simatos

STABILITY PROPERTIES OF
LINEAR FILE-SHARING NETWORKS

Lasse Leskelä				Philippe Robert				Florian Simatos

Lasse Leskelä, Philippe Robert, Florian Simatos: *Stability properties of linear file-sharing networks*; Helsinki University of Technology Institute of Mathematics Research Reports A574 (2009).

Abstract: *File-sharing networks are distributed systems used to disseminate files among a subset of the nodes of the Internet. A file is split into several pieces called chunks, the general simple principle is that once a node of the system has retrieved a chunk, it may become a server for this chunk. A stochastic model is considered for arrival times and durations of time to download chunks. One investigates the maximal arrival rate that such a network can accommodate, i.e., the conditions under which the Markov process describing this network is ergodic. Technical estimates related to the survival of interacting branching processes are key ingredients to establish the stability of these systems. Several cases are considered: networks with one and two chunks where a complete classification is obtained and several cases of a network with n chunks.*

Correspondence

Lasse Leskelä
Helsinki University of Technology
Department of Mathematics and Systems Analysis
PO Box 1100, 02015 TKK
Finland

Philippe Robert, Florian Simatos
INRIA Paris — Rocquencourt
Domaine de Voluceau, BP 105
78153 Le Chesnay
France

lasse.leskela@iki.fi, philippe.robert@inria.fr, florian.simatos@inria.fr

1   Introduction

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

An improved version of this principle consists in splitting the original file into several pieces (called "chunks") so that a given node can retrieve simultaneously several chunks of the same file from different servers. In this case, the rate to get a given file may thus increase significantly. At the same time, the global capacity of the file-sharing system is also increased since a node becomes a server of a chunk as soon as it has retrieved it and not only when it has the whole file. This improvement has interesting algorithmic implications since each node has to establish a matching between chunks and servers. Strategies to maximize the global efficiency of the file sharing systems have to be devised. See for instance Massoulié and Vojnović [12], Bonald *et al.* [4] and Massoulié and Twigg [11].

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

Transient behavior A new file is owned by one node, given there are potentially N other nodes interested by it, how long does it take so that a given node retrieves it ? significant fraction $\alpha \in (0, 1]$ of the N nodes retrieve it ? See Yang and de Veciana [26] and Simatos *et al.* [22]. See also Robert and Simatos [19].

Stationary behavior A constant flow of requests enters, is the capacity of the file-sharing system sufficient to cope with this flow ?

In this paper, the stationary behavior is investigated in a stochastic context: arrival times are random as well as chunk transmission times. In this setting mathematical studies are quite scarce, see Qiu and Srikant [17], Simatos *et al.* [22], Susitaival *et al.* [24] and references therein. A simple strategy to disseminate chunks is considered: chunks are retrieved sequentially and a given node can be the server of only the last chunk it got. See Massoulié and Vojnović [12] and Parvez *et al.* [16] for a detailed motivation of this situation.

In this paper, the sequential scheme for disseminating a file that is divided into n chunks is analyzed. New requests arrive according to a Poisson process at rate λ, and become downloaders of chunk 1. Users who have obtained chunks $1, \ldots, k$ act simultaneously as uploaders of chunk k and downloaders

of chunk $k+1$, and the users who have all the chunks leave the network at rate ν. The transmission rate of chunk k is denoted by μ_k, and x_k is the number of users having obtained chunks $1, \ldots, k$. In this way, the total transmission rate of chunk k in the network is $\mu_k x_k$. The flow of users can be modeled as the linear network depicted in Figure 1.


Figure 1: Transition rates of the linear network outside boundaries.

The main problem analyzed in the paper is the determination of a constant λ^* such that if $\lambda < \lambda^*$ [resp. $\lambda > \lambda^*$], then the associated Markov process is ergodic [resp. transient]. As it will be seen, the constant λ^* may be infinite in some cases so that the file-sharing network is always stable independently of the value of λ. The main technical difficulty to prove stability/instability results for this class of stochastic networks is that, except for the input, the Markov process has unbounded jump rates, in fact proportional to one of the coordinates of the current state. Note that loss networks have also this characteristic but in this case, the stability problem is trivial since the state space is finite. See Kelly [8].

Fluid limits for file-sharing networks. Classically, to analyze the stability properties of stochastic networks, one can use the limits of a scaling of the Markov process, the so-called fluid limits. The scaling consists in speeding up time by the norm $\|x\|$ of the initial state x, by scaling the state vector by $1/\|x\|$ and by letting $\|x\|$ go to infinity. See Bramson [5], Chen and Yao [6] and Robert [18] for example. This scaling is, however, better suited to "locally additive" processes, that is, Markov processes that behave locally as random walks. Since the transition rates are unbounded, it may occur that the corresponding fluid limits have discontinuities; this complicates a lot the analysis of a possible limiting dynamical system. Roughly speaking, this is due to the fact that, because of the unbounded transition rates, events occur on the time scale $t \mapsto t \log \|x\|$ instead of $t \mapsto \|x\| t$. See the case of the $M/M/\infty$ queue in Chapter 9 of Robert [18], and Simatos and Tibi [23] for a discussion of this phenomenon in a related context.

A "fluid scaling" is nevertheless available for file-sharing networks. A possible description for a possible candidate $(x_i(t))$ for this limiting picture would satisfy the following differential equations,

$$
\begin{cases}
\dot{x}_0(t) & = \lambda - \mu_1 x_1(t), \\
\dot{x}_i(t) & = \mu_i x_i(t) - \mu_{i+1} x_{i+1}(t), \quad 1 \leq i \leq n-1, \\
\dot{x}_n(t) & = \mu_n x_n(t) - \nu x_n(t).
\end{cases} \tag{1}
$$

For the sake of simplicity the behavior at the boundaries $\{x : x_i = 0\}$, $i \geq 1$ is ignored in the above equations. This has been, up to now, one of the main

tools to investigate mathematical models of file-sharing networks. See Qiu and Srikant [17], Núñez-Queija and Prabhu [15] for example. In the context of loss networks, an analogous limiting picture can be rigorously justified when the input rates and buffer sizes are scaled by some N and the state variable by $1/N$. This scaling is not useful here, since the problem is precisely of determining the values of λ for which the associated Markov is ergodic whereas in the above scaling λ is scaled. From this point of view Equations (1) are therefore quite informal. They can nevertheless give some insight into the qualitative behavior of these networks but they cannot apparently be used to prove stability results. Their interpretation near boundaries is in particular not clear.

Interacting branching processes. Since scaling techniques do not apply here, one needs to resort to different techniques to study stability: coupling the linear file-sharing network with interacting branching processes is a key idea. For $i \geq 1$, without the departures the process $(X_i(t))$ would be a branching process where individuals give birth to one child at rate μ_i. This description of such a file-sharing system as a branching process is quite natural. It has been used to analyze the transient behavior of these systems. See Yang and de Veciana [26], Dang *et al.* [7] and Simatos *et al.* [22]. A departure for $(X_i(t))$ can be seen as a death of an individual of class i and at the same time as a birth of an individual of class $i+1$. The file-sharing network can thus be described as a system of interacting branching processes with a constant input rate λ.

To tackle the general problem of stability, several key ingredients are used in this paper: Lyapunov functions, coupling arguments and precise estimations of the growth of a branching process killed by another branching process. As it will be seen, several results used come from the branching process formulation of the stochastic model. In particular Section 3 is devoted to the derivation of results concerning killed branching processes. The stability properties of networks with a single-chunk file are analyzed in detail in Section 2. In Section 4, file-sharing networks with n chunks are studied and the case $n = 2$ is investigated thoroughly.

2   Analysis of the single-chunk network

This section is devoted to the study of a class of two-dimensional Markov jump processes $(X_0(t), X_1(t))$, the corresponding Q-matrix Ω_r is given, for

$x = (x_0, x_1) \in \mathbb{N}^2$, by

$$\begin{cases} \Omega_r[(x_0, x_1), (x_0 + 1, x_1)] & = \lambda, \\ \Omega_r[(x_0, x_1), (x_0 - 1, x_1 + 1)] & = \mu r(x_0, x_1)(x_1 \vee 1)\mathbb{1}_{\{x_0 > 0\}}, \\ \Omega_r[(x_0, x_1), (x_0, x_1 - 1)] & = \nu x_1, \end{cases} \qquad (2)$$

where $x \mapsto r(x)$, referred to as the *rate function*, is some fixed function on $\mathbb{N}^2$ with values in $[0, 1]$ and $n \vee m$ denotes $\max(n, m)$ for n, $m \in \mathbb{N}^2$. This corresponds to a more general model than the linear file-sharing network of Figure 1 in the case $n = 1$, where for the sake of simplicity μ_1 is noted μ in this section.

From a modeling perspective, this Markov process describes the following system. Requests for a single file arrive with rate λ, the first component $X_0(t)$ is the number of requests which did not get the file, whereas the second component is the number of requests having the file and acting as servers until they leave the file-sharing network. The constant μ can be viewed as the file transmission rate, and ν as the rate at which servers having all chunks leave. The term $r(x_0, x_1)$ describes the interaction of downloaders and uploaders in the system. The term $x_1 \vee 1$ can be interpreted so that there is one server permanent server in the network, which is contacted if there are no other uploader nodes in the system. A related system where there is always one permanent server for the file can be modeled by replacing the term $x_1 \vee 1$ by $x_1 + 1$. See the remark at the end of this section.

Several related examples of this class of models have been recently investigated. The case

$$r(x_0, x_1) = \frac{x_0}{x_0 + x_1}$$

is considered in Núñez-Queija and Prabhu [15] and Massoulié and Vojnović [12]; in this case the downloading time of the file is neglected. Susitaival *et al.* [24] analyzes the rate function $r(x)$

$$r(x_0, x_1) = 1 \wedge \left(\alpha \frac{x_0}{x_1} \right)$$

with $\alpha > 0$ and $a \wedge b$ denotes $\min(a, b)$ for a, $b \in \mathbb{R}$. This model allows to take into account that a request cannot be served by more than one server. See also Qiu and Srikant [17].

With a slight abuse of notation, for $0 < \delta \leq 1$, the matrix Ω_δ will refer to the case when the function r is identically equal to δ. Note that the boundary condition $x_1 \vee 1$ for departures from the first queue prevents the second coordinate from ending up in the absorbing state 0. Other possibilities are discussed at the end of this section. In the following $(X^r(t)) = (X_0^r(t), X_1^r(t))$ [resp. $(X^\delta(t))$] will denote a Markov process with Q-matrix Ω_r [resp. Ω_δ].

Free process. For $\delta > 0$, Q_δ denotes the following Q-matrix

$$\begin{cases} Q_\delta[(y_0, y_1), (y_0 + 1, y_1)] & = \lambda, \\ Q_\delta[(y_0, y_1), (y_0 - 1, y_1 + 1)] & = \mu\delta(y_1 \vee 1), \\ Q_\delta[(y_0, y_1), (y_0, y_1 - 1)] & = \nu y_1. \end{cases} \qquad (3)$$

The process $(Y^\delta(t)) = (Y_0^\delta(t), Y_1^\delta(t))$, referred to as the free process, will denote a Markov process with Q-matrix Q_δ. Note that the first coordinate Y_0^δ may become negative. The second coordinate $(Y_1^\delta(t))$ of the free process is a classical birth-and-death process. It is easily checked that if ρ_δ defined as $\delta\mu/\nu$ is such that $\rho_\delta < 1$, then $(Y_1^\delta(t))$ is an ergodic Markov process converging in distribution to $Y_1^\delta(\infty)$ and that

$$\lambda^*(\delta) \stackrel{\text{def.}}{=} \nu\mathbb{E}(Y_1^\delta(\infty)) = \mu\mathbb{E}(Y_1^\delta(\infty) \vee 1) = \frac{\delta\mu}{(1 - \rho_\delta)(1 - \log(1 - \rho_\delta))}. \quad (4)$$

When $\rho_\delta > 1$, then the process $(Y^\delta(t))$ converges almost surely to infinity. In the sequel $\lambda^*(1)$ is simply denoted λ^*.

In the following it will be assumed, Condition (C) below, that the rate function r converges to 1 as the first coordinate goes to infinity; as will be seen, the special case $r \equiv 1$ then plays a special role, and so before analyzing the stability properties of $(X^r(t))$, one begins with an informal discussion when the rate function r is identically equal to 1. Since the departure rate from the system is proportional to the number of requests/servers in the second queue, a large number of servers in the second queue gives a high departure rate, irrespectively of the state of the first queue. The input rate of new requests being constant, the real bottleneck with respect to stability is therefore when the first queue is large. The interaction of the two processes $(X_0^1(t))$ and $(X_1^1(t))$ is expressed through the indicator function of the set $\{X_0^1(t) > 0\}$. The second queue $(X_1^1(t))$ locally behaves like the birth-and-death process $(Y_1^1(t))$ as long as $(X_0^1(t))$ is away from 0. The two cases $\rho_1 > 1$ and $\rho_1 < 1$ are considered.

If $\rho_1 > 1$, i.e., $\mu > \nu$, the process $(X_1^1(t))$ is a transient process as long as the first coordinate is non-zero. Consequently, departures from the second queue occur faster and faster. Since, on the other hand, arrivals occur at a steady rate, departures eventually outpace arrivals. The fact that the second queue grows when $(X_0(t))$ is away from 0 stabilizes the system independently of the value of λ, and so the system should be stable for any $\lambda > 0$.

If $\rho_1 < 1$, and as long as $(X_0(t))$ is away from 0, the coordinate $(X_1^1(t))$ locally behaves like the ergodic Markov process $(Y_1^1(t))$. Hence if $(X_0^1(t))$ is non-zero for long enough, the requests in the first queue see in average $\mathbb{E}(Y_1^1(\infty) \vee 1)$ servers which work at rate μ. Therefore, the stability condition for the first queue should be

$$\lambda < \mu\mathbb{E}(Y_1^1(\infty) \vee 1) = \lambda^*$$

where $\lambda^* = \lambda^*(1)$ is defined by Equation (4). Otherwise if $\lambda > \lambda^*$, the system should be unstable.

Markovian notations. In the following, one will use the following convention, if $(U(t))$ is a Markov process, the index u of $\mathbb{P}_u((U(t)) \in \cdot)$ will refer to the initial condition of this Markov process.

Transience and recurrence criteria for $(X^r(t))$.

Proposition 2.1 (Coupling). *If $X^r(0) = Y^1(0) \in \mathbb{N}^2$, there exists a coupling of the processes $(X^r(t))$ and $(Y^1(t))$ such that the relation*

$$X_0^r(t) \geq Y_0^1(t) \text{ and } X_1^r(t) \leq Y_1^1(t), \tag{5}$$

holds for all $t \geq 0$ and for any sample path.
  For any $0 \leq \delta \leq 1$, if

$$\tau_\delta = \inf\{t \geq 0 : r(X^r(t)) \leq \delta\} \text{ and } \sigma = \inf\{t \geq 0 : X_0^r(t) = 0\},$$

and if $X^1(0) = Y^\delta(0) \in \mathbb{N}^2$ then there exists a coupling of the processes $(X^r(t))$ and $(Y^\delta(t))$ such that, for any sample path, the relation

$$X_0^r(t) \leq Y_0^\delta(t) \text{ and } X_1^r(t) \geq Y_1^\delta(t) \tag{6}$$

holds for all $t \leq \tau_\delta \wedge \sigma$.

Proof. Let $X^r(0) = (x_0, x_1)$ and $Y^1(0) = (y_0, y_1)$ be such that $x_0 \geq y_0$ and $x_1 \leq y_1$, one has to prove that the processes $(X^r(t))$ and $(Y^1(t))$ can be constructed such that Relation (5) holds at the time of the next jump of one of them. See Leskelä [10] for the existence of couplings using analytical, nonconstructive techniques.
  The arrival rates in the first queue are the same for both processes. If $x_1 < y_1$, a departure from the second queue for $(Y^1(t))$ or $(X^r(t))$ preserves the order relation (5) and if $x_1 = y_1$, this departure occurs at the same rate for both processes and thus the corresponding instant can be chosen at the same (exponential) time. For the departures from the first to the second queue, the departure rate for $(X^r(t))$ is $\mu r(x_0, x_1)(x_1 \vee 1)\mathbb{1}_{\{x_0 > 0\}} \leq \mu(y_1 \vee 1)$ which is the departure rate for $(Y^1(t))$, hence the corresponding departure instants can be taken in the reverse order so that Relation (5) also holds at the next jump instant. The first part of the proposition is proved.
  The rest of the proof is done in a similar way: The initial states $X^r(0) = (x_0, x_1)$ and $Y^\delta(0) = (y_0, y_1)$ are such that $x_0 \leq y_0$ and $x_1 \geq y_1$. With the killing of the processes at time $\tau_\delta \wedge \sigma$ one can assume additionally that $x_0 \neq 0$ and that the relation $r(x_0, x_1) \geq \delta$ holds; Under these assumptions one can check by inspecting the next transition that (6) holds. The proposition is proved. $\square$

Proposition 2.2. *Under the condition $\mu < \nu$, the relation*

$$\liminf_{t \to +\infty} \frac{X_0^r(t)}{t} \geq \lambda - \lambda^*$$

holds almost surely. In particular, if $\mu < \nu$ and $\lambda > \lambda^*$, then the process $(X^r(t))$ is transient.

8

Proof. By Proposition 2.1, one can assume that there exists a version of $(Y^1(t))$ such that $X_0^r(0) = Y_0^1(0)$ and the relation $X_0^r(t) \geq Y_0^1(t)$ holds for any $t \geq 0$. From Definition (3) of the Q-matrix of $(Y^1(t))$, one has, for $t \geq 0$,

$$Y^1(t) = Y^1(0) + \mathcal{N}_\lambda(t) - A(t),$$

where $(\mathcal{N}_\lambda(t))$ is a Poisson process with parameter λ and $(A(t))$ is the number of arrivals (jumps of size 1) for the second coordinate $(Y_1^1(t))$: in particular

$$\mathbb{E}(A(t)) = \mu \mathbb{E}\left(\int_0^t Y_1^1(s) \vee 1\, ds\right).$$

Since $(Y_1^1(t))$ is an ergodic Markov process under the condition $\mu < \nu$, the ergodic theorem in this setting gives that

$$\lim_{t \to +\infty} \frac{1}{t} A(t) = \lim_{t \to +\infty} \frac{1}{t} \mathbb{E}(A(t)) = \mu \mathbb{E}\left(Y_1^1(\infty) \vee 1\right) = \lambda^*,$$

by Equation (4), hence $(Y_0^1(t)/t)$ converges almost surely to $\lambda - \lambda^*$. The proposition is proved. $\qquad\square$

The next result establishes the ergodicity result of this section.

Proposition 2.3. *If the rate function r is such that, for any $x_1 \in \mathbb{N}$,*

$$\lim_{x_0 \to +\infty} r(x_0, x_1) = 1, \tag{C}$$

and if $\mu \geq \nu$, or if $\mu < \nu$ and $\lambda < \lambda^*$ with

$$\lambda^* = \frac{\mu}{(1-\rho)(1 - \log(1-\rho))}, \tag{7}$$

and $\rho = \mu/\nu$, then $(X^r(t))$ is an ergodic Markov process.

Note that Condition (C) is satisfied for the functions r considered in the models considered by Núñez-Queija and Prabhu [15] and in Susitaival *et al.* [24]. See above.

Proof. If $x = (x_0, x_1) \in \mathbb{R}^2$, $|x|$ denotes the norm of x, $|x| = |x_0| + |x_1|$. The proof uses Foster's criterion as stated in Robert [18, Theorem 9.7]. If there exist constants K_0, K_1, t_0, t_1 and $\eta > 0$ such that, for $x = (x_0, x_1) \in \mathbb{N}^2$,

$$\mathbb{E}_{(x_0,x_1)}(|X^r(t_1)| - |x|) \leq -t_1, \text{ if } x_1 \geq K_1, \tag{8}$$
$$\mathbb{E}_{(x_0,x_1)}(|X^r(t_0)| - |x|) \leq -\eta t_0, \text{ if } x_0 \geq K_0 \text{ and } x_1 < K_1, \tag{9}$$

then the Markov process $(X^r(t))$ is ergodic.

Relation (8) is straightforward to establish: if $x_1 \geq K_1$, one gets, by considering only K_1 of the x_1 initial servers in the second queue and the Poisson arrivals, that

$$\mathbb{E}_{(x_0,x_1)}(|X^r(1)| - |x|) \leq \lambda - K_1(1 - e^{-\nu}),$$

hence it is enough to take $t_1 = 1$ and $K_1 = (\lambda + 1)/(1 - e^{-\nu})$ to have Relation (8).

One has therefore to establish Inequality (9). Let τ_δ and σ be the stopping times introduced in Proposition 2.1, one first proves an intermediate result: for any $t > 0$ and any $x_1 \in \mathbb{N}$,

$$\lim_{x_0 \to +\infty} \mathbb{P}_{(x_0,x_1)}(\sigma \wedge \tau_\delta \leq t) = 0. \tag{10}$$

Fix $x_1 \in \mathbb{N}$ and $t \geq 0$: for $\varepsilon > 0$, there exists D_1 such that

$$\mathbb{P}_{x_1}\left(\sup_{0 \leq s \leq t} Y_1^1(s) \geq D_1\right) \leq \varepsilon,$$

from Proposition 2.1, this gives the relation valid for all $x_0 \geq 0$,

$$\mathbb{P}_{(x_0,x_1)}\left(\sup_{0 \leq s \leq t} X_1^r(s) \geq D_1\right) \leq \varepsilon.$$

By Condition (C), there exists $\gamma \geq 0$ (that depends on x_1) such that $r(x_0, x_1) \geq \delta$ when $x_0 \geq \gamma$. As long as $(X^r(t))$ stays in the subset $\{(y_0, y_1) : y_1 \leq D_1\}$, the transition rates of the first component $(X_0^r(t))$ are uniformly bounded. Consequently, there exists K such that, for $x_0 \geq K$,

$$\mathbb{P}_{(x_0,x_1)}\left[\sup_{s \leq t} X_0^r(s) \leq \gamma, \ \sup_{s \leq t} X_1^r(s) \leq D_1,\right] \leq \varepsilon.$$

Relation (10) follows from the last two inequalities and the identity

$$\mathbb{P}_{(x_0,x_1)}(\sigma \wedge \tau_\delta \leq t) \leq \mathbb{P}_{(x_0,x_1)}\left(\sup_{s \leq t} X_0^r(s) \leq \gamma\right).$$

One returns to the proof of Inequality (9). By definition of the Q-matrix of the process $(X^r(t))$,

$$\mathbb{E}_{(x_0,x_1)}(|X^r(t)| - |x|) = \lambda t - \nu \int_0^t \mathbb{E}_{(x_0,x_1)}(X_1^r(u))du, \ x \in \mathbb{N}^2, \ t \geq 0.$$

For any $x \in \mathbb{N}^2$, there exists a version of $(Y^\delta(t))$ with initial condition $Y^\delta(0) = X^r(0) = x$, and such that Relation (6) holds for $t < \tau_\delta \wedge \sigma$, in particular

$$\mathbb{E}_x(X_1^r(t)) \geq \mathbb{E}_x(X_1^r(t); t < \tau_\delta \wedge \sigma)$$
$$\geq \mathbb{E}_x(Y_1^\delta(t); t < \tau_\delta \wedge \sigma) = \mathbb{E}_x(Y_1^\delta(t)) - \mathbb{E}_x(Y_1^\delta(t); t \geq \tau_\delta \wedge \sigma).$$

Cauchy-Schwarz inequality shows that for any $t \geq 0$ and $x \in \mathbb{N}^2$

$$\int_0^t \mathbb{E}_x(Y_1^\delta(u); \tau_\delta \wedge \sigma \leq u)\,du \leq \int_0^t \sqrt{\mathbb{E}_x\left[\left(Y_1^\delta(u)\right)^2\right]}\sqrt{\mathbb{P}_x(\tau_\delta \wedge \sigma \leq u)}\,du$$
$$\leq \sqrt{\mathbb{P}_x(\tau_\delta \wedge \sigma \leq t)} \int_0^t \sqrt{\mathbb{E}_x\left[\left(Y_1^\delta(u)\right)^2\right]}\,du,$$

10

by gathering these inequalities, and by using the fact that the process $(Y_1^\delta(t))$ depends only on x_1 and not x_0, one finally gets the relation

$$\frac{1}{t}\mathbb{E}_x(|X(t)| - |x|) \le \lambda - \frac{\nu}{t}\int_0^t \mathbb{E}_{x_1}(Y_1^\delta(u))\,du + c(x_1,t)\sqrt{\mathbb{P}_x(\tau_\delta \wedge \sigma \le t)} \quad (11)$$

with

$$c(x_1,t) = \frac{\nu}{t}\int_0^t \sqrt{\mathbb{E}_{x_1}\left[\left(Y_1^\delta(u)\right)^2\right]}\,du.$$

Two cases are considered:

1. If $\mu > \nu$, if $\delta < 1$ is such that $\delta\mu > \nu$, the process $(Y_1^\delta(t))$ is transient, so that

$$\lim_{t\to+\infty}\frac{1}{t}\int_0^t \mathbb{E}_{x_1}(Y_1^\delta(u))\,du = +\infty,$$

   for each $x_1 \ge 0$.

2. If $\mu < \nu$, one takes $\delta = 1$, or if $\mu = \nu$, one takes $\delta < 1$ close enough to 1 so that $\lambda < \lambda^*(\delta)$. In both cases, $\lambda < \lambda^*(\delta)$ and the process $(Y_1^\delta(t))$ converges in distribution, hence

$$\lim_{t\to+\infty}\frac{1}{t}\int_0^t \mathbb{E}_{x_1}(Y_1^\delta(u))\,du = \nu\mathbb{E}\left(Y_1^\delta(\infty)\right) = \lambda^*(\delta) > \lambda$$

   for each $x_1 \ge 0$.

Consequently in both cases, there exist constants $\eta > 0$, $\delta < 1$ and $t_0 > 0$ such that for any $x_1 \le K_1$,

$$\lambda - \nu\frac{1}{t_0}\int_0^{t_0} \mathbb{E}_{x_1}(Y_1^\delta(u))du \le -\eta, \tag{12}$$

with Relation (11), one gets that if $x_1 \le K_1$ then

$$\frac{1}{t_0}\mathbb{E}_x(|X(t_0)| - |x|) \le -\eta + c^*\sqrt{\mathbb{P}_x(\tau_\delta \wedge \sigma \le t_0)},$$

where $c^* = \max(c(n,t_0), 0 \le n \le K_1)$. By Identity (10), there exists K_0 such that, for all $x_0 \ge K_0$ and $x_1 \le K_1$, the relation

$$c^*\sqrt{\mathbb{P}_{(x_0,x_1)}(\tau_\delta \wedge \sigma \le t_0)} \le \frac{\eta}{2}$$

holds. This relation and the inequalities (12) and (11) give Inequality (9). The proposition is proved. $\qquad\square$

Another boundary condition. The boundary condition $x_1 \vee 1$ in the transition rates of $(X(t))$, Equation (2), prevents the second coordinate from ending up in the absorbing state 0. It amounts to suppose that a permanent server gets activated when no node may offer the file. Another way to avoid this absorbing state is to suppose that a permanent node is always active, which gives transition rates with $x_1 + 1$ instead. This choice was for instance made in Núñez-Queija and Prabhu [15]. All our results apply for this other boundary condition: the only difference that is when $\nu > \mu$, the value of the threshold λ^* of Equation (4) is given by the quantity $\lambda^* = \mu\nu/(\nu - \mu)$.

3 Yule processes with deletions

This section introduces the tools which are necessary in order to generalize the results of the previous section to the multi-chunk case $n \geq 2$. A Yule process $(Y(t))$ with rate $\mu > 0$ is a Markovian branching process with Q-matrix

$$q_Y(x, x + 1) = \mu x, \quad \forall x \geq 0. \tag{13}$$

An individual gives birth to a child, or equivalently splits into two particles, with rate μ. Let (σ_n) be the split times of a Yule process started with one particle, it is not difficult to check that, for $n \geq 1$,

$$\sigma_n \overset{\text{dist.}}{=} \sum_{\ell=1}^{n} \frac{E_\ell^\mu}{\ell} \overset{\text{dist.}}{=} \max(E_1^\mu, \ldots, E_n^\mu),$$

where (E_ℓ^μ) are i.i.d. exponential random variables with parameter μ. If $\lambda > \mu$ then, by using Fubini's Theorem,

$$\mathbb{E}\left(\sum_{\ell=1}^{+\infty} e^{-\lambda \sigma_\ell}\right) = \mathbb{E}\left(\sum_{\ell=1}^{+\infty} \int_0^{+\infty} \lambda e^{-\lambda x} \mathbb{1}_{\{\sigma_\ell \leq x\}} \, dx\right) = \int_0^{+\infty} \lambda e^{-\lambda x} \sum_{\ell=1}^{+\infty} \mathbb{P}(\sigma_\ell \leq x) \, dx$$

$$= \int_0^{+\infty} \lambda e^{-\lambda x} \frac{1 - e^{-\mu x}}{e^{-\mu x}} \, dx = \frac{\mu}{\lambda - \mu} < +\infty. \tag{14}$$

In this section one considers some specific results on variants of this stochastic model when some individuals are killed. In terms of branching processes, this amounts to prune the tree, i.e., to cut some edges of the tree, and the subtree attached to it. This procedure is fairly common for branching processes, in the Crump-Mode-Jagers model for example, see Kingman [9]. See also Neveu [14] or Aldous and Pitman [1]. Two situations are considered: the first one when the deletions are part of the internal dynamics, so that each individual dies out after an exponential time, and the other when killings are given by an exogenous process and occur at fixed (random or deterministic) epochs.

Constant death rate and regeneration. Let $(Z(t))$ be the birth-and-death process whose Q-matrix Q_Z is given by, for $\mu_Z > 0$ and $\nu > 0$,

$$q_Z(z, z + 1) = \mu_Z(z \vee 1) \quad \text{and} \quad q_Z(z, z - 1) = \nu z. \tag{15}$$

The lifetime of an individual is exponentially distributed with parameter ν, and the process restarts with one individual after some time when it hits 0. This process can be described equivalently as a time-changed $M/M/1$ queue or as a sequence of independent branching processes. As it will be seen these two viewpoints are complementary.

In the rest of this part, μ_Z and ν are fixed, $(Z(t))$ is the Markov process with Q-matrix Q_Z, (σ_n) is the sequence of times of its positive jumps, the birth instants, and $(B_\sigma(t))$ is the corresponding counting process of (σ_n), for $t \geq 0$,

$$B_\sigma(t) = \sum_{i \geq 1} \mathbb{1}_{\{\sigma_i \leq t\}}.$$

Proposition 3.1 (Queueing Representation)**.** *If $Z(0) = z \in \mathbb{N}$, then*

$$(Z(t), t \geq 0) \stackrel{dist.}{=} (L(C(t)), t \geq 0), \tag{16}$$

where $(L(t))$ is the process of the number of jobs of an $M/M/1$ queue with input rate μ_Z and service rate ν and with $L(0) = z$ and $C(t) = \inf\{s > 0 : A(s) > t\}$, where

$$A(t) = \int_0^t \frac{1}{1 \vee L(u)} \, du.$$

Proof. It is not difficult to check that the process $(M(t)) \stackrel{def.}{=} (L(C(t)))$ has the Markov property. Let Q_M be its Q-matrix. For $z \geq 0$,

$$\mathbb{P}(L(C(h)) = z + 1 \mid L(0) = z) = \mu_Z \mathbb{E}(C(h)) + o(h) = \mu_Z(z \vee 1)h + o(h),$$

hence $q_M(z, z+1) = \mu_Z(z \vee 1)$. Similarly $q_M(z, z-1) = \nu z$. The proposition is proved. $\qquad\square$

Corollary 3.1. *For any $\gamma > (\mu_Z - \nu) \vee 0$ and $z = Z(0) \in \mathbb{N}$,*

$$\mathbb{E}_z \left(\sum_{n=1}^{+\infty} e^{-\gamma \sigma_n} \right) < +\infty. \tag{17}$$

Proof. Proposition 3.1 shows that, in particular, the sequences of positive jumps of $(Z(t))$ and of $(L(C(t)))$ have the same distribution. Hence, if $\mathcal{N}_{\mu_Z} = (t_n)$ is the arrival process of the $M/M/1$ queue, a Poisson process with parameter μ_Z, then, with the notations of the above proposition, the relation

$$(\sigma_n) \stackrel{dist.}{=} (A(t_n))$$

holds. By using standard martingale properties of stochastic integrals with respect to Poisson processes, see Rogers and Williams [20], one gets for $t \geq 0$,

$$\mathbb{E}_z \left(\sum_{n \geq 1} e^{-\gamma A(t_n)} \right) = \mathbb{E}_z \left(\int_0^\infty e^{-\gamma A(s)} \mathcal{N}_{\mu_Z}(ds) \right) = \mu_Z \mathbb{E}_z \left(\int_0^\infty e^{-\gamma A(s)} \, ds \right)$$

$$= \mu_Z \int_0^\infty e^{-\gamma u} \mathbb{E}_z \left(Z(u) \vee 1 \right) \, du, \tag{18}$$

where Relation (16) has been used for the last equality. Kolmogorov's equation for the process $(Z(t))$ gives that

$$\phi(t) \stackrel{def.}{=} \mathbb{E}_z(Z(t)) = \mu_Z \int_0^t \mathbb{E}_z \left(Z(u) \vee 1 \right) \, du - \nu \int_0^t \mathbb{E}_z \left(Z(u) \right) \, du$$

$$\leq (\mu_Z - \nu) \int_0^t \phi(u) \, du + \mu_Z t,$$

therefore, by Gronwall's Lemma,

$$\phi(t) \leq \phi(0) + \mu_Z \int_0^t u e^{(\mu_Z - \nu)u} \, du \leq z + \frac{\mu_Z}{\mu_Z - \nu} t e^{(\mu_Z - \nu)t}.$$

From Equation (18), one concludes that

$$\mathbb{E}_z\left(\sum_n e^{-\gamma\sigma_n}\right) = \mathbb{E}_z\left(\sum_n e^{-\gamma A(t_n)}\right) < +\infty.$$

The proposition is proved. $\qquad\square$

A branching process. Before hitting 0, the Markov process $(Z(t))$ whose Q-matrix is given by Relation (15) can be seen a Bellman-Harris branching process. Its Malthusian parameter is given by $\alpha = \mu_Z - \nu$. See Athreya and Ney [3]. In this setting, it describes the evolution of a population of independent particles, at rate $\lambda \stackrel{\text{def.}}{=} \mu_Z + \nu$ each of these particles either splits into two particles with probability $p \stackrel{\text{def.}}{=} \mu_Z/(\mu_Z+\nu)$ or dies. These processes will be referred to as (p,λ)-branching processes in the sequel.

A (p,λ)-branching process survives with positive probability only when $p > 1/2$, in which case the probability of extinction q is equal to $q = (1 - p)/p = \nu/\mu_Z$. The main (and only) difference with a branching process is that Z regenerates after hitting 0. When it regenerates, it again behaves as a (p,λ)-branching process (started with one particle), until it hits 0 again.

Proposition 3.2 (Branching Representation). *If $Z(0) = z \in \mathbb{N}$ and $(\widetilde{Z}(t))$ is a (p,λ)-branching process started with $z \in \mathbb{N}$ particles and $\widetilde{T}$ its extinction time, then*

$$(Z(t), 0 \le t \le T) \stackrel{dist.}{=} (\widetilde{Z}(t), 0 \le t \le \widetilde{T}),$$

where $T = \inf\{t \ge 0 : Z(t) = 0\}$ is the hitting time of 0 by $(Z(t))$.

Corollary 3.2. *Suppose that $\mu_Z > \nu$. Then $\mathbb{P}_z$-almost surely for any $z \ge 0$, there exists a finite random variable $Z(\infty)$ such that,*

$$\lim_{t\to+\infty} e^{-(\mu_Z-\nu)t} Z(t) = Z(\infty) \quad and \quad Z(\infty) > 0.$$

Proof. When $\mu_Z > \nu$, the process $(Z(t))$ couples in finite time with a supercritical (p,λ)-branching process $(\widetilde{Z}(t))$ conditioned on non-extinction; this follows readily from Proposition 3.2 (or see the Appendix for details). Since for any supercritical (p,λ)-branching process, $(\exp(-(\mu_Z - \nu)t)\widetilde{Z}(t))$ converges almost surely to a finite random variable $\widetilde{Z}(\infty)$, positive on the event of non-extinction (see Nerman [13]), one gets the desired result. $\qquad\square$

Due to its technicality, the proof of the following result is postponed to the Appendix; this result is used in the proof of Proposition 3.5.

Proposition 3.3. *Suppose that $\mu_Z > \nu$, if*

$$\eta^*(x) = \frac{2 - x - \sqrt{x(4 - 3x)}}{2(1 - x)}, \quad 0 < x < 1, \tag{19}$$

then for any $0 < \eta < \eta^*(\nu/\mu_Z)$,

$$\sup_{z\ge0}\left[\mathbb{E}_z\left(\sup_{t\ge\sigma_1}\left(e^{\eta(\mu_Z-\nu)t}B_\sigma(t)^{-\eta}\right)\right)\right] < +\infty.$$

14

A Yule process killed at fixed instants. In this part, it is assumed that, provided that it is non-empty, at epochs σ_n, $n \geq 1$, an individual is removed from the population of an ordinary Yule process $(Y(t))$ with rate μ_W starting with $Y(0) = w \in \mathbb{N}$ individuals. It is assumed that (σ_n) is some fixed non-decreasing sequence. It will be shown that the process $(W(t))$ obtained by killing one individual of $Y(t)$ at each of the successive instants (σ_n) survives with positive probability when the series with general term $(\exp(-\mu_W \sigma_n))$ converges.

In the following, a related result will be considered in the case where the sequence (σ_n) is given by the sequence of birth times of the process $(Z(t))$ introduced above. See Alsmeyer [2] and the references therein for related models.

One denotes

$$\kappa = \inf\{n \geq 1 : W(\sigma_n) = 0\}.$$

The process $(W(t))$ can be represented in the following way

$$W(t) = Y(t) - \sum_{i=1}^{\kappa} X_i(t) \mathbb{1}_{\{\sigma_i \leq t\}}, \tag{20}$$

where, for $1 \leq i \leq \kappa$ and $t \geq \sigma_i$, $X_i(t)$ is the total number of children at time t in the original Yule process of the ith individual killed at time σ_i. In terms of trees, $(W(t))$ can be seen as a subtree of $(Y(t))$: for $1 \leq i \leq \kappa$, $(X_i(t))$ is the subtree of $(Y(t))$ associated with the ith particle killed at time σ_i.

It is easily checked that $(X_i(t - \sigma_i), t \geq \sigma_i)$ is a Yule process starting with one individual and, since a killed individual cannot have one of his descendants killed, that the processes

$$(\widetilde{X}_i(t)) = (X_i(t + \sigma_i), t \geq 0), \quad 1 \leq i \leq \kappa,$$

are independent Yule processes.

For any process $(U(t))$, one denotes

$$(M_U(t)) \overset{\text{def.}}{=} \left(e^{-\mu_W t} U(t)\right). \tag{21}$$

If $(\widetilde{X}(t))$ is a Yule process with rate μ_W, the martingale $(M_{\widetilde{X}}(t))$ converges almost surely and in L_2 to a random variable $M_{\widetilde{X}}(\infty)$ with an exponential distribution with mean $\widetilde{X}(0)$, and by Doob's Inequality

$$\mathbb{E}\left(\sup_{t \geq 0} M_{\widetilde{X}}(t)^2\right) \leq 2 \sup_{t \geq 0} \mathbb{E}\left(M_{\widetilde{X}}(t)^2\right) < +\infty.$$

See Athreya and Ney [3]. Consequently

$$e^{-\mu_W t} W(t) = M_Y(t) - \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(t - \sigma_i) \mathbb{1}_{\{\sigma_i \leq t\}},$$

and for any $t \geq 0$,

$$\sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(t - \sigma_i) \mathbb{1}_{\{\sigma_i \leq t\}} \leq \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} \sup_{s \geq 0} M_{\widetilde{X}_i}(s).$$

Assume now that $\sum_{i \geq 1} e^{-\mu_W \sigma_i} < +\infty$: then the last expression is integrable, and Lebesgue's Theorem implies that $(M_W(t)) = (\exp(-\mu_W t)W(t))$ converges almost surely and in L_2 to

$$M_W(\infty) = M_Y(\infty) - \sum_{i=1}^{\kappa} e^{-\mu_W \sigma_i} M_{\widetilde{X}_i}(\infty).$$

Clearly, for some w^* large enough and then for any $w \geq w^*$, one has

$$\mathbb{E}_w(M_W(\infty)) \geq w - \sum_{i=1}^{+\infty} e^{-\mu_W \sigma_i} > 0,$$

in particular $\mathbb{P}_w(M_W(\infty) > 0) > 0$ and $\mathbb{P}_w(W(t) \geq 1, \forall t \geq 0) > 0$. If $Y(0) = w < w^*$ and $\sigma_1 > 0$, then $\mathbb{P}_w(Y(\sigma_1) \geq w^* + 1) > 0$ and therefore, by translation at time σ_1, the same conclusion holds when the sequence $(\exp(-\mu_W \sigma_i))$ has a finite sum. The following proposition has thus been proved.

Proposition 3.4. *Let $(W(t))$ be a process growing as a Yule process with rate μ_W and for which individuals are killed at non-decreasing instants (σ_n) with $\sigma_1 > 0$. If*

$$\sum_{i=1}^{+\infty} e^{-\mu_W \sigma_i} < +\infty,$$

then as t gets large, and for any $w \geq 1$, the variable $(\exp(-\mu_W t)W(t))$ converges $\mathbb{P}_w$-almost surely and in L_2 to a finite random variable $M_W(\infty)$ such that $\mathbb{P}_w(M_W(\infty) > 0) > 0$.

The previous proposition establishes the minimal results needed in Section 4. However, Kolmogorov's Three-Series, see Williams [25], can be used in conjunction with Fatou's Lemma to show that $(W(t))$ dies out almost surely when the series with general term $(\exp(-\mu_W \sigma_n))$ diverges.

A Yule process killed at the birth instants of a Bellman-Harris process. In this subsection, one considers a Yule process $(Y(t))$ with parameter μ_W with Q-matrix defined by Relation (13) and an independent Markov process $(Z(t))$ with Q-matrix defined by Relation (15). In particular $\mu_Z - \nu$ is the Malthusian parameter of $(Z(t))$. A process $(W(t))$ is defined by killing one individual of $(Y(t))$ at each of the birth instants (σ_n) of $(Z(t))$. As before $(B_\sigma(t))$ denotes the counting process association to the non-decreasing sequence (σ_n),

$$B_\sigma(t) = \sum_{i \geq 1} \mathbb{1}_{\{\sigma_i \leq t\}}.$$

Proposition 3.5. *Assume that $\mu_Z - \nu > \mu_W$, and let H_0 be the extinction time of $(W(t))$, i.e.,*

$$H_0 = \inf\{t \geq 0 : W(t) = 0\},$$

then the random variable H_0 is almost surely finite and:

(i) $Z(H_0) - Z(0) \leq e^{\mu_W H_0} M_Y^*$ *where*

$$M_Y^* = \sup_{t \geq 0} e^{-\mu_W t} Y(t).$$

(ii) There exists a finite constant C such that for any $z \geq 0$ and $w \geq 1$,

$$\mathbb{E}_{(w,z)}(H_0) \leq C\left(\log(w) + 1\right). \tag{22}$$

Note that the subscript (w, z) refers to the initial state of the Markov process $(W(t), Z(t))$.

Proof. Define $\alpha = \mu_Z - \nu$. Concerning the almost sure finiteness of H_0, note that Equation (20) entails that $W(t) \leq Y(t) - B_\sigma(t)$ for all $t \geq 0$ on the event $\{H_0 = +\infty\}$. As t goes to infinity, both $\exp(-\mu_W t)Y(t)$ and $\exp(-\alpha t)B_\sigma(t)$ converge almost surely to positive and finite random variables (see Nerman [13]), which implies, when $\alpha = \mu_Z - \nu > \mu_W$, that $W(t)$ converges to $-\infty$ on $\{H_0 = +\infty\}$, and so this event is necessarily of probability zero.

The first point (i) of the proposition comes from Identity (20) at $t = H_0$:

$$Z(H_0) - Z(0) \leq B_\sigma(H_0) \leq Y(H_0) \leq e^{\mu_W H_0} M_Y^*. \tag{23}$$

By using the relation $\exp(x) \geq x$, Equation (22) follows from the following bound: for any $\eta < \eta^*(\nu/\mu_Z)$ (recall that η^* is given by Equation (19)),

$$\sup_{w \geq 1, z \geq 0} \left[w^{-\eta} \mathbb{E}_{(w,z)}\left(e^{\eta(\alpha - \mu_W)H_0}\right)\right] < +\infty. \tag{24}$$

So all is left to prove is this bound. Under $\mathbb{P}_{(w,z)}$, $(Y(t))$ can be represented as the sum of w i.i.d. Yule processes, and so $M_Y^* \leq M_{Y,1}^* + \cdots + M_{Y,w}^*$ with $(M_{Y,i}^*)$ i.i.d. distributed like M_Y^* under $\mathbb{P}_{(1,z)}$; Inequality (23) then entails that

$$e^{(\alpha - \mu_W)H_0} \leq \left(\sum_{i=1}^{w} M_{Y,i}^*\right) \times \sup_{t \geq \sigma_1}\left(e^{\alpha t}/B_\sigma(t)\right).$$

By independence of $(M_{Y,i}^*)$ and $(B_\sigma(t))$, Jensen's inequality gives for any $\eta < 1$

$$\mathbb{E}_{(w,z)}\left(e^{\eta(\alpha - \mu_W)H_0}\right) \leq w^\eta \left(\mathbb{E}\left(M_{Y,1}^*\right)\right)^\eta \mathbb{E}_z\left(\sup_{t \geq \sigma_1}\left(e^{\eta\alpha t}B_\sigma(t)^{-\eta}\right)\right),$$

hence the bound (24) follows from Proposition 3.3. $\qquad\square$

One concludes this section with a Markov chain which will be used in Section 4. Define recursively the sequence (V_n) by, $V_0 = v$ and

$$V_{n+1} = \sum_{k=1}^{A_n(V_n)} I_k, \quad n \geq 0, \tag{25}$$

where (I_k) are identically distributed integer valued random variables independent of V_n and $A_n(V_n)$, and such that $\mathbb{E}(I_1) = p$ for some $p \in (0, 1)$. For $v > 0$, $A_n(v)$ is an independent random variable with the same distribution as $Z(H_0)$ under $\mathbb{P}_{(1,v)}$, i.e., with the initial condition $(W(0), Z(0)) = (1, v)$.

The above equation (25) can be interpreted as a branching process with immigration, see Seneta [21], or also as an autoregressive model.

Proposition 3.6. *Under the condition $\mu_Z - \nu > \mu_W$, if (V_n) is the Markov chain defined by Equation (25) and, for $K \geq 0$,*

$$N_K = \inf\{n \geq 0 : V_n \leq K\},$$

then there exist $\gamma > 0$ and $K \in \mathbb{N}$ such that

$$\mathbb{E}(N_K | V_0 = v) \leq \frac{1}{\gamma} \log(1 + v), \quad \forall v \geq 0. \tag{26}$$

The Markov chain (V_n) is in particular positive recurrent.

Proof. For $V_0 = v \in \mathbb{N}$, Jensen's Inequality and Definition (25) give the relation

$$\mathbb{E}_v \log\left(\frac{1 + V_1}{1 + v}\right) \leq \mathbb{E}_{(1,v)} \log\left[\frac{1 + pZ(H_0)}{1 + v}\right]. \tag{27}$$

From Proposition 3.5 and by using the same notations, one gets that, under $\mathbb{P}_{(1,v)}$,

$$Z(H_0) \leq v + e^{\mu_W H_0} M_Y^*,$$

where $(Y(t))$ is a Yule process starting with one individual. By looking at the birth instants of $(Z(t))$, it is easily checked that the random variable H_0 under $\mathbb{P}_{(1,v)}$ is stochastically bounded by H_0 under $\mathbb{P}_{(1,0)}$. The integrability of H_0 under $\mathbb{P}_{(1,0)}$ (proved in Proposition 3.5) and of M_Y^* give that the expression

$$\log\left(\frac{1 + p(v + e^{\mu_W H_0} M_Y^*)}{1 + v}\right)$$

bounding the right hand side of Relation (27) is also an integrable random variable under $\mathbb{P}_{(1,0)}$. Lebesgue's Theorem gives therefore that

$$\limsup_{v \to +\infty} \left[\mathbb{E}_v \log\left(\frac{1 + V_1}{1 + v}\right)\right] \leq \log p < 0.$$

Consequently, one concludes that $v \mapsto \log(1 + v)$ is a Lyapunov function for the Markov chain (V_n), i.e., if $\gamma = -(\log p)/2$, there exists K such that for $v \geq K$,

$$\mathbb{E}_v \log(1 + V_1) - \log(1 + v) \leq -\gamma.$$

Foster's criterion, see Theorem 8.6 of Robert [18], implies that (V_n) is indeed ergodic and that Relation (26) holds. $\square$

4   Analysis of the multi-chunk network

In this section it is assumed that a file of n chunks is distributed by the file-sharing network within the following framework, corresponding to Figure 1. Chunks are delivered in the sequential order, and, for $k \geq 1$, requests with chunks $1, \ldots, k$ provide service for requests with one less chunk.

For $0 \leq k < n$ and $t \geq 0$, the variable $X_k(t)$ denotes the number of requests downloading the $(k+1)$st chunk; for $k = n$, $X_n(t)$ is the number of requests having all the chunks. When taking into account the boundaries in the transition rates described in Figure 1, one gets the following Q-matrix for the $(n+1)$-dimensional Markov process $(X_k(t), 0 \leq k \leq n)$:

$$Q(f)(x) = \lambda[f(x+e_0)-f(x)]+\sum_{k=1}^{n}\mu_k(x_k\vee 1)[f(x+e_k-e_{k-1})-f(x)]\mathbb{1}_{\{x_{k-1}>0\}}$$
$$+ \nu x_n[f(x-e_n)-f(x)],$$

where $x \in \mathbb{N}^{n+1}$, $f : \mathbb{N}^{n+1} \to \mathbb{R}_+$ is a function and for, $0 \leq k \leq n$, $e_k \in \mathbb{N}^{n+1}$ is the kth unit vector. Note that, as before, to avoid absorbing states, it is assumed that there is a server for the kth chunk when $x_k = 0$. The first section corresponds to the case $n = 2$ in a more general setting.

It is first shown in Proposition 4.1 that the network is stable for sufficiently small input rate λ. Proposition 4.2 studies the analog of the two-dimensional case with $\mu > \nu$, i.e., when $\mu_1 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$, it is proved that the network is stable for any input rate λ. When this condition fails, it is shown that for $n = 2$ the network can only accommodate a finite input rate.

Proposition 4.1. *Under the condition*

$$\sum_{k=1}^{n}\frac{\lambda}{\mu_k} < 1, \tag{28}$$

the Markov process $(X(t))$ is ergodic for any $\nu > 0$.

Condition (28) is obviously not sharp as can be seen in the case $n = 1$ analyzed in Section 2. But the proposition shows that there is always a positive threshold λ^* such that the system is stable when $\lambda < \lambda^*$.

Proof. For $x \in \mathbb{N}^{n+1}$ and $(\alpha_k) \in \mathbb{R}^{n+1}$, define $f(x) = \alpha_0 x_0 + \cdots + \alpha_n x_n$, then

$$Q(f)(x) = \lambda\alpha_0 - \sum_{k=1}^{n}(\alpha_{k-1} - \alpha_k)\mu_k(x_k \vee 1)\mathbb{1}_{\{x_{k-1}>0\}} - \nu x_n\alpha_n.$$

For $\varepsilon > 0$, one can choose (α_k) so that $\alpha_0 = 1$ and

$$\alpha_{k-1} - \alpha_k = \frac{\lambda}{\mu_k} + \varepsilon, \quad 1 \leq k \leq n,$$

hence

$$\alpha_n = 1 - \left(n\varepsilon + \sum_{i=1}^{n} \frac{\lambda}{\mu_k} \right),$$

so that, for ε small enough, the α_k's, $0 \le k \le n$ are decreasing and positive under the condition of the proposition; in particular the set $\{x : f(x) \le K\}$ is finite for any $K \ge 0$.

Take $K = (1 + \lambda)/\nu$, then if $x \in \mathbb{N}^{n+1}$ is such that $f(x) \ge K$, either $x_k > 0$ for some $0 \le k \le n-1$ and in this case

$$Q(f)(x) \le \lambda - \mu_{k+1}(\alpha_k - \alpha_{k+1}) = -\varepsilon\mu_{k+1} < 0,$$

or $x_n \ge K$ so that

$$Q(f)(x) \le \lambda - \nu K = -1 < 0.$$

A Lyapunov function criteria for Markov processes shows that this implies that the Markov process $(X(t))$ is ergodic. See Proposition 8.14 of Robert [18] for example. $\qquad\square$

Decreasing service rates. The analog of the "good" case $\mu > \nu$ is proved in the next proposition.

Proposition 4.2. *Under the condition $\mu_1 > \mu_2 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$, the Markov process $(X(t)) = (X_k(t), 0 \le k \le n)$ describing the linear file-sharing network is ergodic for any $\lambda \ge 0$.*

Proof. The proof proceeds in two steps: first coupling arguments with Yule processes allow to prove (30); then one can use the same technique as in the proof of Proposition 2.3, see Robert [18, Theorem 9.7].

Step 1 (coupling). Let $(W_n(t))$ be the process with Q-matrix defined by Relation (15) with $\mu_Z = \mu_n$ and starting at $W_n(0) = w_n \ge 1$. Since $\mu_n > \nu$, the process $(\exp(-(\mu_n - \nu)t)W_n(t))$ converges almost surely to a finite and positive random variable $M_{W_n}(\infty)$ by Corollary 3.2. Moreover, since $\mu_{n-1} > \mu_n - \nu > 0$, Corollary 3.1 entails that the birth instants (σ_ℓ^n) of this process are such that

$$\sum_{\ell \ge 1} e^{-\mu_{n-1}\sigma_\ell^n} < +\infty, \quad \text{almost surely.}$$

Let $(Y_{n-1}(t))$ be an independent Yule process with parameter μ_{n-1} with initial condition $Y_{n-1}(0) = w_{n-1} \ge 1$ and $(W_{n-1}(t))$ the resulting process when its individuals are killed at the instants (σ_ℓ^n) of births of $(W_n(t))$: the previous equation and Proposition 3.4 show that $(W_{n-1}(t))$ can survive forever with a positive probability.

Let $(Y_{n-2}(t))$ be an independent Yule process starting from $w_{n-2} \ge 1$ with parameter μ_{n-2}. Define $(W_{n-2}(t))$ the resulting process when the individuals of $(Y_{n-2}(t))$ are killed at the birth instants (σ_ℓ^{n-1}) of $(W_{n-1}(t))$. Since $\mu_{n-2} > \mu_{n-1}$, the birth instants $(\widetilde{\sigma}_\ell^{n-1})$ of $(Y_{n-1}(t))$ satisfy

$$\sum_{\ell=1}^{+\infty} e^{-\mu_{n-2}\widetilde{\sigma}_\ell^{n-1}} < +\infty$$

almost surely by Equation (14) (which still holds for a Yule process starting with more than one particle). Since the birth instants (σ_ℓ^{n-1}) of $(W_{n-1}(t))$ are a subsequence of $(\widetilde{\sigma}_\ell^{n-1})$, the same relationship holds for (σ_ℓ^{n-1}), and therefore, with a positive probability, the three processes $(e^{-(\mu_n-\nu)t}W_n(t))$, $(e^{-\mu_{n-1}t}W_{n-1}(t))$ and $(e^{-\mu_{n-2}t}W_{n-2}(t))$ converge simultaneously to positive and finite random variables $M_{W_n}(\infty)$, $M_{W_{n-1}}(\infty)$ and $M_{W_{n-2}}(\infty)$, respectively. This construction can be repeated inductively to give the existence of n processes $(W_k(t), k = 1, \ldots, n)$ such that (σ_ℓ^k) is the sequence of birth times of W_k, W_n is the birth-and-death process with Q-matrix (15), W_k for $1 \le k \le n-1$ is a Yule process with parameter μ_k killed at (σ_ℓ^{k+1}), and the event $\mathcal{E} = \{M_{W_1}(\infty) > 0, \ldots, M_{W_n}(\infty) > 0\}$ has a positive probability. On this event, $W_k(t) \ge 1$ for all $t \ge 0$ and $1 \le k \le n-1$, and

$$\lim_{t \to +\infty} W_n(t) = +\infty.$$

For $0 \le k \le n-1$, one defines $(X_k^S(t)) = (X_{k,n-k}^S(t), \ldots, X_{k,n}^S(t))$, the kth saturated system, as the $(k+1)$-dimensional Markov process with generator

$$Q_k^S(f)(x) = \mu_{n-k}(x_{n-k} \vee 1)[f(x + e_{n-k}) - f(x)]$$
$$+ \sum_{\ell=1}^{k} \mu_{n-k+\ell}(x_{n-k+\ell} \vee 1)[f(x + e_{n-k+\ell} - e_{n-k+\ell-1}) - f(x)]\mathbb{1}_{\{x_{n-k+\ell-1}>0\}}$$
$$+ \nu x_n[f(x - e_n) - f(x)], \quad (29)$$

where $x \in \mathbb{N}^{k+1}$ and $f : \mathbb{N}^{k+1} \to \mathbb{R}_+$ is an arbitrary function. Compared with the process $(X_\ell(t), 1 \le \ell \le n)$ with generator Q, it amounts to look at the $k+1$ last queues $(X_{n-k}(t), \ldots, X_n(t))$ under the assumption that the queue $n-k-1$ is saturated, i.e., $X_{n-k-1}(t) \equiv +\infty$ for all $t \ge 0$.

Note that for any $0 \le k \le n-1$, the transition rates of the Markov processes $(W_{n-\ell}(t), 0 \le \ell \le k)$ and $(X_{k,n-\ell}^S(t), 0 \le \ell \le k)$ are identical as long as no coordinate hits 0; one thus concludes that, with positive probability, the relation

$$\lim_{t \to +\infty} X_{k,n}^S(t) = +\infty$$

holds when $X_{k,n-\ell}^S(0) \ge 1$, $\ell = 0, \ldots, k$. Consequently, since the set $(\mathbb{N} - \{0\})^{k+1}$ can be reached with positive probability from any initial state in $\mathbb{N}^{k+1}$ by $(X_k^S(t))$, then

$$\lim_{t \to +\infty} \mathbb{E}(X_{k,n}^S(t)) = +\infty. \quad (30)$$

Step 2 (Foster's criterion). We use Foster's criterion as stated in Theorem 9.7 of Robert [18]. First we inspect the case when $X_n(0)$ is large, then the case when $X_n(0)$ is bounded and $X_{n-1}(0)$ is large, etc. . . The key idea is that when $X_{n-k-1}(0)$ is large, then the process $(X_{n-k}(t), \ldots, X_n(t))$ essentially behaves as the process $(X_k^S(t))$, for which Relation (30) ensures that the output rate is arbitrarily large.

Let $X(0) = x = (x_k) \in \mathbb{N}^{n+1}$, since the last queue serves at rate ν each request, for $t \geq 0$,

$$\mathbb{E}(\|X(t)\|) \leq \|x\| + \lambda t - x_n \left(1 - e^{-\nu t}\right),$$

where $\|x\| = x_0 + \cdots + x_n$ for $x = (x_0, \ldots, x_n) \in \mathbb{N}^{n+1}$. Define $t_n = 1$ and let K_n be such that $\lambda t_n - K_1(1 - \exp(-\nu)) \leq -1$, so that the relation

$$\mathbb{E}_x(\|X(t_n)\|) - \|x\| \leq -1,$$

holds when $x_n \geq K_n$.

From Equation (30) with $k = 0$, one gets that there exists some t_{n-1} such that for any $x_n \leq K_n$,

$$\nu \int_0^{t_{n-1}} \mathbb{E}_{x_n} \left(X_{0,n}^S(u)\right) du \geq \lambda t_{n-1} + 2.$$

The two processes $(X_0^S(t))$ and $(X(t))$ can be built on the same probability space such that if they start from the same initial state, then the two processes $(X_{0,n}^S(t))$ and $(X_n(t))$ are identical as long as $X_{n-1}(t)$ stays positive. Since moreover the hitting time $\inf\{t \geq 0 : X_{n-1}(t) = 0\}$ goes to infinity as x_{n-1} goes to infinity for any $x_n \leq K_n$, one gets that there exists K_{n-1} such that if $x_{n-1} \geq K_{n-1}$ and $x_n < K_n$, then the relation

$$\mathbb{E}_x(\|X(t_{n-1})\|) - \|x\| = \lambda t_{n-1} - \nu \int_0^{t_{n-1}} \mathbb{E}_x(X_n(u)) \, du$$
$$\leq \lambda t_{n-1} - \left(\nu \int_0^{t_{n-1}} \mathbb{E}_{x_n} \left(X_{0,n}^S(u)\right) du - 1\right) \leq -1$$

holds.

By induction, one gets in a similar way that there exist constants $t_n, \ldots, t_0$ and $K_n, \ldots, K_0$ such that for any $0 \leq \ell \leq n$, if $x_n \leq K_n$, $x_{n-1} \leq K_{n-1}$, $\ldots$, $x_{n-\ell+1} \leq K_{n-\ell+1}$ and $x_{n-\ell} > K_{n-\ell}$, then

$$\mathbb{E}_x(\|X(t_{n-\ell})\|) - \|x\| \leq -1.$$

Theorem 8.13 of Robert [18] shows that $(X(t))$ is an ergodic Markov process. The proposition is proved. $\qquad \square$

Analysis of the two-chunk network. In this subsection, one investigates the case when the monotonicity condition $\mu_1 > \cdots > \mu_{n-1} > \mu_n - \nu > 0$ fails. In general we conjecture the existence of bottlenecks which implies that the network can only accommodate a finite input rate. For instance, when $\mu_n - \nu < 0$, then it is easily seen that the network is unstable for $\lambda > \lambda^*$ where λ^* is defined in Equation (32) below.

The first non-trivial case occurs for $n = 2$, for which the monotonicity condition breaks in two situations, either when $\mu_2 - \nu > \mu_1$ or when $\mu_2 < \nu$. The latter case can be dealt in fact with the exact same arguments as before.

See Proposition 4.4. The actual difficulty is when $\mu_2 - \nu > \mu_1$: then the stationary behavior of $(X_2(t))$ is linked to the stationary behavior of the first saturated model $(X_1^S(t))$ defined through its Q-matrix (29). The difficulty in this case is that one needs to compare two processes which grow exponentially fast.

Proposition 4.3. *Assume that $\mu_2 - \nu > \mu_1$, then the first saturated process $(X_1^S(t))$ with Q-matrix defined by Equation (29) is ergodic.*

Corollary 4.1. *If $\mu_2 - \nu > \mu_1$ and if*

$$\lambda_2^* \stackrel{def.}{=} \nu \mathbb{E}_{\pi^S} \left(X_{1,2}^S(0) \right),$$

where π^S is the invariant distribution of the Markov process $(X_1^S(t))$, then the process $(X(t)) = (X_k(t), k = 0, 1, 2)$ describing the linear file-sharing network with parameters λ, μ_1, μ_2 and ν is ergodic for $\lambda < \lambda_2^*$ and transient for $\lambda > \lambda_2^*$.

Sketch of Proof. The proof of the transience when $\lambda > \lambda_2^*$ follows similarly as in Section 2: when $X_0(0)$ is large, the process $(X_1(t), X_2(t))$ can be coupled for some time with the second saturated system $(X_1^S(t))$. Since the output rate λ_2^* of this system is smaller than the input rate λ, this implies that $(X_0(t))$ builds up, and it can indeed be shown that $X_0(t)/t$ converges almost surely to $\lambda - \lambda_2^*$.

The ergodicity when $\lambda < \lambda_2^*$ is slightly more complicated, but it involves the same arguments as the ones employed in the proof of Proposition 4.2. The details are omitted. □

Proof of Proposition 4.3. Denote $(X_1^S(t)) = (X_{1,1}^S(t), X_{1,2}^S(t))$, then as long as the first coordinate $X_{1,1}^S$ is positive, the process $(X_1^S(t))$ has the same distribution as $(W(t), Z(t))$ introduced in Section 3: $(Z(t))$ is a Bellman-Harris process with Malthusian parameter $\mu_2 - \nu$ and $(W(t))$ is a Yule process with parameter μ_1 killed at times of births of $(Z(t))$.

By Proposition 3.5 and since $\mu_2 - \nu > \mu_1$, one has that $(X_{1,1}^S(t))$ returns infinitely often to 0. When $(X_{1,1}^S(t))$ is at 0 it jumps to 1 after an exponential time with parameter μ_1, one denotes by $(E_{\mu_1,n})$ the corresponding i.i.d. sequence of successive residence times at 0. One defines the sequence (S_n) by induction, $S_0 = 0$ and then

$$S_{n+1} = \inf\{t > S_n : X_{1,1}^S(t) = 0\} + E_{\mu_1,n+1}, \ n \geq 0.$$

For $n \geq 1$, $X_{1,1}^S(S_n) = 1$ and for $n \geq 0$, define $M_n \stackrel{def.}{=} X_{1,2}^S(S_n)$. With the notations of Proposition 3.5, $(X_{1,1}^S(t))$ hits 0 after a duration of $H_{0,n}$ and at that time $(X_{1,2}^S(t))$ is at $Z(H_{0,n})$ with the initial condition $Z(0) = M_n$; while $X_{1,1}^S$ is still at 0, the dynamics of $X_{1,2}^S$ is simple, since it just empties. Finally, at time $S_{n+1} = S_n + H_{0,n} + E_{\mu_1,n+1}$, $(X_{1,1}^S(t))$ returns to 1 and at this instant the location of $(X_{1,2}^S(t))$ is given by

$$X_{1,2}^S(S_{n+1}) = M_{n+1} = \sum_{i=1}^{Z(H_{0,n})} \mathbb{1}_{\{E_{\nu,i} > E_{\mu_1,n+1}\}},$$

23

where $(E_{\nu,i})$ are i.i.d. exponential random variables with parameter ν, the ith variable being the residence time of the ith request in node 2. Consequently, $(M_n, n \geq 1)$ is a Markov chain whose transitions are defined by Relation (25) with $p = \nu/(\nu + \mu_1)$; note that $(M_n, n \geq 0)$ has the same dynamics only when $X_{1,1}^S(0) = 1$.

Define for any $K > 0$ the stopping time T_K

$$T_K = \inf\{t \geq 0 : X_{1,2}^S(t) \leq K, X_{1,1}^S(t) = 1\}.$$

The ergodicity of $(X_1^S(t))$ will follow from the finiteness of $\mathbb{E}_{(x_1,x_2)}(T_K)$ for some K large enough and for arbitrary $x = (x_1, x_2) \in \mathbb{N}^2$. The strong Markov property of $(X_1^S(t))$ applied at time S_1 gives

$$\mathbb{E}_{(x_1,x_2)}(T_K) \leq 2\mathbb{E}_{(x_1,x_2)}(S_1) + \mathbb{E}_{(x_1,x_2)}\left[\mathbb{E}_{(1,X_{1,2}^S(S_1))}(T_K)\right],$$

and so one only needs to study T_K conditioned on $\{X_{1,1}^S(0) = 1\}$ since $\mathbb{E}_{(x_1,x_2)}(S_1)$ is finite in view of Proposition 3.5.

Then, on this event and with N_K defined in Proposition 3.6, the identity

$$T_K = \sum_{i=0}^{N_K}(H_{0,i} + E_{\mu_1,i}) \tag{31}$$

holds. For $i \geq 0$, the Markov property of $(M_n, n \geq 0)$ gives

$$\mathbb{E}_{(x_1,x_2)}\left(H_{0,i}\mathbb{1}_{\{i \leq N_K\}}\right) = \mathbb{E}_{(x_1,x_2)}\left(\mathbb{E}_{(1,M_i)}(H_0)\,\mathbb{1}_{\{i \leq N_K\}}\right)$$

With the same argument as in the proof of Proposition 3.6, one has

$$\mathbb{E}_{(1,M_i)}(H_0) \leq \mathbb{E}_{(1,0)}(H_0) < +\infty,$$

with Equations (31) and (26) of Proposition (3.6), one gets that for some $\gamma > 0$ and some $K > 0$,

$$\mathbb{E}_{(x_1,x_2)}(T_K) \leq 2\mathbb{E}_{(x_1,x_2)}(S_1) + C\left(1 + \mathbb{E}_{(x_1,x_2)}\left[\log\left(1 + X_{1,2}^S(S_1)\right)\right]\right)$$

with the constant $C = (\mathbb{E}_{(1,0)}(H_0) + 1/\mu_2)/\gamma$. This last term is finite for any (x_1, x_2) in view of Proposition 3.5, which proves the proposition. $\square$

Proposition 4.4. *If $\nu > \mu_2$ and*

$$\lambda^* \stackrel{def.}{=} \frac{\mu_2}{(1 - \mu_2/\nu)(1 - \log(1 - \mu_2/\nu))}, \tag{32}$$

then the Markov process $(X(t)) = (X_k(t), k = 0, 1, 2)$ is transient if $\lambda > \lambda^*$ and ergodic if $\lambda < \lambda^*$.

Sketch of Proof. The result for transience comes directly from the fact that the last coordinate is stochastically dominated by the birth-and-death process $(Y_1^1(t))$ of Section 2.

As before, the arguments employed in the proof of Proposition 4.2 to prove ergodicity can also be used, for this reason they are only sketched. One has in fact to consider the following situations:

— If there are many customers in the last queue, then the total number of customers instantaneously decreases.

— If there are many customers in the second queue, then the last queue has time to get close to stationarity, the input rate is λ and the output rate is λ^*.

— Finally, if there are many customers in the first queue, then it is easily seen that the second queue builds up, since it grows like a Yule process killed at times (σ_n) where the sequence (σ_n) essentially grows linearly since the last queue is stable. Hence the second queue reaches high values and the last queue offers an output rate of λ^*.

Hence when $\lambda < \lambda^*$, the Markov process $(X(t))$ is ergodic. $\qquad\square$

A   Proof of Proposition 3.3

In this appendix the notations of Section 3 are used. Since the random variable $(B_\sigma(t) \mid Z(0) = 0)$ is stochastically smaller than $(B_\sigma(t) \mid Z(0) = z)$ for any $z \in \mathbb{N}$, it is enough to show that for $\eta < \eta^*(\nu/\mu_Z)$

$$\mathbb{E}_0 \left[\sup_{t \geq \sigma_1} \left(e^{\eta\alpha t} B_\sigma(t)^{-\eta} \right) \right] < +\infty,$$

where $\alpha = \mu_Z - \nu > 0$.

Note that the process $(B_\sigma(t + \sigma_1), t \geq 0)$ under $\mathbb{P}_0$ has the same distribution as $(B_\sigma(t) + 1, t \geq 0)$ under $\mathbb{P}_1$, and by independence of σ_1, an exponentially random variable with parameter μ_Z, and $(B_\sigma(t + \sigma_1), t \geq 0)$, one gets

$$\mathbb{E}_0 \left[\sup_{t \geq \sigma_1} \left(e^{\eta\alpha t} B_\sigma(t)^{-\eta} \right) \right] = \mathbb{E}_0 \left(e^{\eta\alpha\sigma_1} \right) \mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta\alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \right].$$

Since $\alpha < \mu_Z$ and $\eta^*(\nu/\mu_Z) < 1$, then $\mathbb{E}_0 \left(\exp(\eta\alpha\sigma_1) \right)$ is finite, and all one needs to prove is that the second term is finite as well.

Define τ as the last time $Z(t) = 0$:

$$\tau = \sup\{t \geq 0 : Z(t) = 0\},$$

with the convention that $\tau = +\infty$ if $(Z(t))$ never returns to 0. Recall that, because of the assumption $\mu_Z > \nu$, with probability 1, the process $(Z(t))$ returns to 0 a finite number of times.

Conditioned on the event $\{\tau = +\infty\}$, the process $(Z(t))$ is a (p, λ)-branching process conditioned on survival, with $\lambda = \mu_Z + \nu$ and $p = \mu_Z/\lambda$. Such a branching process conditioned on survival can be decomposed as $Z = Z_{(1)} + Y$, where $(Y(t))$ is a Yule process $(Y(t))$ with parameter α. See Athreya and Ney [3]. Consequently, for any $0 < \eta < 1$,

$$\mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta\alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \mid \tau = +\infty \right] \leq \mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta\alpha t} Y(t)^{-\eta} \right) \right].$$

Since the nth split time t_n of $(Y(t))$ is distributed like the maximum of n i.i.d. exponential random variables, $Y(t)$ for $t \geq 0$ is geometrically distributed with parameter $1 - e^{-\alpha t}$, hence,

$$\sup_{t \geq 0} \left[e^{\eta \alpha t} \mathbb{E}_1 \left(\frac{1}{Y(t)^\eta} \right) \right] = \sup_{t \geq 0} \left[e^{-(1-\eta)\alpha t} \sum_{k \geq 1} \frac{(1 - e^{-\alpha t})^{k-1}}{k^\eta} \right]$$

$$\leq \sup_{0 \leq u \leq 1} \left[(1 - u)^{1-\eta} \sum_{k \geq 1} \frac{u^{k-1}}{k^\eta} \right].$$

For $0 < u < 1$, the relation

$$(1 - u)^{1-\eta} \sum_{k \geq 1} \frac{u^{k-1}}{k^\eta} \leq (1 - u)^{1-\eta} \int_0^\infty \frac{u^x}{(1 + x)^\eta} \, dx,$$

$$= \left(\frac{1 - u}{-\log u} \right)^{1-\eta} \int_0^\infty \frac{e^{-x}}{(x - \log u)^\eta} \, dx,$$

holds, hence

$$\sup_{t \geq 0} \left[e^{\eta \alpha t} \mathbb{E}_1 \left(\frac{1}{Y(t)^\eta} \right) \right] < +\infty.$$

The process $(e^{-\alpha t} Y(t))$ being a martingale, by convexity the process $(e^{\eta \alpha t} Y(t)^{-\eta})$ is a non-negative sub-martingale. For any $\eta \in (0, 1)$ Doob's L_p inequality gives the existence of a finite $q(\eta) > 0$ such that

$$\mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta \alpha t} Y(t)^{-\eta} \right) \right] \leq q(\eta) \sup_{t \geq 0} \left[e^{\eta \alpha t} \mathbb{E}_1 \left(\frac{1}{Y(t)^\eta} \right) \right] < +\infty.$$

The following result has therefore been proved.

Lemma A.1. *For any* $0 < \eta < 1$,

$$\mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \middle| \tau = +\infty \right] < +\infty.$$

On the event $\{\tau < +\infty\}$, $(Z(t))$ hits a geometric number of times 0 and then couples with a (p, λ)-branching process conditioned on survival. On this event,

$$\sup_{t \geq 0} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right)$$

$$= \max \left(\sup_{0 \leq t \leq \tau} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right), \sup_{t \geq \tau} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \right)$$

$$\leq e^{\eta \alpha \tau} \left(1 + \sup_{t \geq 0} \left(e^{\eta \alpha t} \left(B_\sigma'(t) + 1 \right)^{-\eta} \right) \right)$$

26

where $B'_\sigma(t)$ for $t \geq \tau$ is the number of births in $(\tau, t]$ of a (p, λ)-branching process conditioned on survival and independent of the variable τ, consequently

$$\mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \middle| \tau < +\infty \right] \leq \mathbb{E}_1 \left(e^{\eta \alpha \tau} \middle| \tau < +\infty \right)$$
$$\times \left(1 + \mathbb{E}_1 \left[\sup_{t \geq 0} \left(e^{\eta \alpha t} \left(B_\sigma(t) + 1 \right)^{-\eta} \right) \middle| \tau = +\infty \right] \right).$$

In view of Lemma A.1, the proof of Proposition 3.3 will be finished if one can prove that

$$\mathbb{E}_1 \left(e^{\eta \alpha \tau} \middle| \tau < +\infty \right) < +\infty,$$

which actually comes from the following decomposition: under $\mathbb{P}_1(\cdot \mid \tau < +\infty)$, the random variable τ can be written as

$$\tau = \sum_{k=1}^{1+G} (T_k + E_{\mu_Z, k})$$

where G is a geometric random variable with parameter $q = \nu / \mu_Z$, (T_k) is an i.i.d. sequence with the same distribution as the extinction time of a (p, λ)-branching process starting with one particle and conditioned on extinction and $(E_{\mu_Z, k})$ are i.i.d. exponential random variables with parameter μ_Z.

Since q is the probability of extinction of a (p, λ)-branching process started with one particle, $G + 1$ represents the number of times $(Z(t))$ hits 0 before going to infinity. This representation entails

$$\mathbb{E}_1 \left(e^{\eta \alpha \tau} \mid \tau < +\infty \right) = \mathbb{E} \left(\gamma(\eta)^{G+1} \right) \quad \text{where} \quad \gamma(\eta) = \mathbb{E} \left(e^{\eta \alpha (T_1 + E_{\mu_Z, 1})} \right).$$

A (p, λ)-branching process conditioned on extinction is actually a $(1 - p, \lambda)$-branching process. See again Athreya and Ney [3]. Thus T_1 satisfies the following recursive distributional equation:

$$T_1 \overset{\text{dist.}}{=} E_\lambda + \mathbb{1}_{\{\xi = 2\}} (T_1 \vee T_2),$$

where $\mathbb{P}(\xi = 2) = 1 - p$ and E_λ is an exponential random variable with parameter λ. This equation yields

$$\mathbb{P}(T_1 \geq t) \leq e^{-\lambda t} + 2\lambda (1 - p) \int_0^t \mathbb{P}(T_1 \geq t - u) e^{-\lambda u} \, du,$$

and Gronwall's Lemma applied to the function $t \mapsto \exp(\lambda t) \mathbb{P}(T_1 \geq t)$ gives that

$$\mathbb{P}(T_1 \geq t) \leq e^{(\lambda - 2\lambda p) t} = e^{(\nu - \mu_Z) t}$$

hence for any $0 < \eta < 1$,

$$\mathbb{E}_1 \left(e^{\eta \alpha T_1} \right) \leq \frac{1}{1 - \eta}.$$

Since G is a geometric random variable with parameter q, $\mathbb{E}\left(\gamma(\eta)^G\right)$ is finite if and only if $\gamma(\eta) < q$. Since finally

$$\gamma(\eta) = \frac{\mu_Z}{\mu_Z - \eta\alpha}\mathbb{E}\left(e^{\eta\alpha T_1}\right) \leq \frac{\mu_Z}{(1-\eta)(\mu_Z - \eta\alpha)},$$

one can easily check that $\gamma(\eta) < q$ for $\eta < \eta^*(\nu/\mu_Z)$ as defined by Equation (19), which concludes the proof of Proposition 3.3.

References

[1] David Aldous and Jim Pitman, *Tree-valued Markov chains derived from Galton-Watson processes*, Annales de l'Institut Henri Poincaré. Probabilités et Statistiques **34** (1998), no. 5, 637–686.

[2] Gerold Alsmeyer, *On the Galton-Watson Predator-Prey Process*, Annals of Applied Probability **3** (1993), no. 1, 198–211.

[3] K. B. Athreya and P. E. Ney, *Branching processes*, Springer, 1972.

[4] Thomas Bonald, Laurent Massoulié, Fabien Mathieu, Diego Perino, and Andrew Twigg, *Epidemic live streaming: optimal performance trade-offs*, Proceedings of SIGMETRICS'08 (New York, NY, USA), ACM, 2008, pp. 325–336.

[5] Maury Bramson, *Stability of queueing networks*, Lecture Notes in Mathematics, vol. 1950, Springer, Berlin, 2008, Lectures from the 36th Probability Summer School held in Saint-Flour, July 2–15, 2006.

[6] Hong Chen and David D. Yao, *Fundamentals of queueing networks*, Springer-Verlag, New York, 2001, Performance, asymptotics, and optimization, Stochastic Modelling and Applied Probability.

[7] T. D. Dang, R. Pereczes, and S. Molnár, *Modeling the population of file-sharing peer-to-peer networks with branching processes*, IEEE Symposium on Computers and Communications (ISCC'07) (Aveiro, Portugal), July 2007.

[8] F.P. Kelly, *Loss networks*, Annals of Applied Probability **1** (1991), no. 3, 319–378.

[9] J. F. C. Kingman, *The first birth problem for an age-dependent branching process.*, Annals of Probability **3** (1975), no. 5, 790–801.

[10] L. Leskelä, *Stochastic relations of random variables and processes*, J. Theor. Probab. (2009), To appear.

[11] Laurent Massoulié and Andrew Twigg, *Rate-optimal schemes for Peer-to-Peer live streaming*, Performance Evaluations **65** (2008), no. 11-12, 804–822.

[12] Laurent Massoulié and Milan Vojnović, *Coupon replication systems*, Proceedings of SIGMETRICS'05 (Banff, Alberta, Canada), no. 1, June 2005, pp. 2–13.

[13] Olle Nerman, *On the convergence of supercritical general (C-M-J) branching processes*, Z. Wahrscheinlichkeitstheorie verw. Gebiete **57** (1981), 365–395.

[14] J. Neveu, *Erasing a branching tree*, Advances in Applied Probability (1986), no. suppl., 101–108.

[15] R. Núñez-Queija and B. J. Prabhu, *Scaling laws for file dissemination in P2P networks with random contacts*, Proceedings of IWQoS, 2008.

[16] Nadim Parvez, Carey Williamson, Anirban Mahanti, and Niklas Carlsson, *Analysis of bittorrent-like protocols for on-demand stored media streaming*, SIGMETRICS '08: Proceedings of the 2008 ACM SIGMETRICS international conference on Measurement and modeling of computer systems (New York, NY, USA), ACM, 2008, pp. 301–312.

[17] Dongyu Qiu and R. Srikant, *Modeling and performance analysis of bittorrent-like peer-to-peer networks*, SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications (New York, NY, USA), ACM, 2004, pp. 367–378.

[18] Philippe Robert, *Stochastic networks and queues*, Stochastic Modelling and Applied Probability Series, vol. 52, Springer, New-York, June 2003.

[19] Philippe Robert and Florian Simatos, *Occupancy schemes associated to Yule processes*, Advances in Applied Probability **41** (2009), no. 2, To Appear.

[20] L. C. G. Rogers and David Williams, *Diffusions, Markov processes, and martingales. Vol. 2: Itô calculus*, John Wiley & Sons Inc., New York, 1987.

[21] E. Seneta, *On the supercritical branching process with immigration*, Mathematical Biosciences **7** (1970), 9–14.

[22] Florian Simatos, Philippe Robert, and Fabrice Guillemin, *A queueing system for modeling a file sharing principle*, Proceedings of SIGMETRICS'08 (New York, NY, USA), ACM, 2008, pp. 181–192.

[23] Florian Simatos and Danielle Tibi, *Spatial homogenization in a stochastic network with mobility*, Annals of Applied Probability (2009), To Appear.

[24] Riikka Susitaival, Samuli Aalto, and Jorma Virtamo, *Analyzing the dynamics and resource usage of P2P file sharing by a spatio-temporal model*, International Workshop on P2P for High Performance Computational Sciences, 2006.

[25] David Williams, *Probability with martingales*, Cambridge University Press, 1991.

[26] Xiangying Yang and Gustavo de Veciana, *Service capacity of peer to peer networks*, Proceedings of IEEE Infocom'04, ACM, 2004, pp. 2242–2252.

A568   Olavi Nevanlinna
       Computing the spectrum and representing the resolvent
       April 2009

A567   Antti Hannukainen, Sergey Korotov, Michal Krizek
       On a bisection algorithm that produces conforming locally refined simplicial
       meshes
       April 2009

A566   Mika Juntunen, Rolf Stenberg
       A residual based a posteriori estimator for the reaction–diffusion problem
       February 2009

A565   Ehsan Azmoodeh, Yulia Mishura, Esko Valkeila
       On hedging European options in geometric fractional Brownian motion market
       model
       February 2009

A564   Antti H. Niemi
       Best bilinear shell element: flat, twisted or curved?
       February 2009

A563   Dmitri Kuzmin, Sergey Korotov
       Goal-oriented a posteriori error estimates for transport problems
       February 2009

A562   Antti H. Niemi
       A bilinear shell element based on a refined shallow shell model
       December 2008

A561   Antti Hannukainen, Sergey Korotov, Michal Krizek
       On nodal superconvergence in 3D by averaging piecewise linear, bilinear, and
       trilinear FE approximations
       December 2008

A560   Sampsa Pursiainen
       Computational methods in electromagnetic biomedical inverse problems
       November 2008

A573   Mika Juntunen
        Finite element methods for parameter dependent problems
        June 2009

A572   Bogdan Bojarski
        Differentiation of measurable functions and Whitney–Luzin type structure
        theorems
        June 2009

A571   Lasse Leskelä
        Computational methods for stochastic relations and Markovian couplings
        June 2009

A570   Janos Karatson, Sergey Korotov
        Discrete maximum principles for FEM solutions of nonlinear elliptic systems
        May 2009

A569   Antti Hannukainen, Mika Juntunen, Rolf Stenberg
        Computations with finite element methods for the Brinkman problem
        April 2009