AIMD ALGORITHMS AND EXPONENTIAL FUNCTIONALS

BY FABRICE GUILLEMIN, PHILIPPE ROBERT[1] AND BERT ZWART

France Telecom, *INRIA and Eindhoven University of Technology*

The behavior of a connection transmitting packets into a network according to a general additive-increase multiplicative-decrease (AIMD) algorithm is investigated. It is assumed that loss of packets occurs in clumps. When a packet is lost, a certain number of subsequent packets are also lost (correlated losses). The stationary behavior of this algorithm is analyzed when the rate of occurrence of clumps becomes arbitrarily small. From a probabilistic point of view, it is shown that exponential functionals associated to compound Poisson processes play a key role. A formula for the fractional moments and some density functions are derived. Analytically, to get the explicit expression of the distributions involved, the natural framework of this study turns out to be the *q*-calculus. Different loss models are then compared using concave ordering. Quite surprisingly, it is shown that, for a fixed loss rate, the correlated loss model has a higher throughput than an uncorrelated loss model.

1. Introduction. TCP (Transmission Control Protocol) is the main data transmission protocol of the Internet. It is designed to adapt to the various traffic conditions of the present network: a TCP connection between a source and a destination progressively increases its transmission rate until it receives some indication that the capacity along its path in the network is almost fully utilized. On the other hand, when the capacity of the network cannot accommodate the traffic (when delays and timeouts affect the connection), the data rate of the connection is drastically reduced. More specifically, a given connection has a variable W which gives the maximum number of packets that can be transmitted without receiving any acknowledgement from the destination. The variable W is called the *congestion window size*. If all the $W + 1$ packets are successfully transmitted, then W is increased by 1 (progressive test of the available capacity of the network), so that W packets can be sent for the next round. Otherwise W is divided by 2 (detection of congestion). TCP uses an additive-increase multiplicative-decrease (AIMD) algorithm with additive increment 1 and multiplicative decay $\delta = 1/2$. An AIMD algorithm can be described as follows:

$$W \to \begin{cases} W + 1, & \text{if no loss among the } W \text{ packets,} \\ \lfloor \delta W \rfloor, & \text{otherwise,} \end{cases}$$

where $\lfloor x \rfloor$ is the integer part of $x \in \mathbb{R}$.

The emergence of TCP as the ubiquitous data transmission protocol has motivated over the past ten years a huge amount of research for modeling a TCP connection experiencing packet loss. Since the initial work by Floyd [12], who derived via simulation an asymptotic estimate for the throughput of a TCP connection experiencing a constant loss rate α (the $c/\sqrt{\alpha}$ formula), several studies have refined the results obtained by Floyd. The paper by Padhye, Firoiu, Towsley and Kurose [21] gives an asymptotic estimate for the throughput of a TCP connection experiencing independent losses of packets. This result has been obtained via an approximation of a finite state Markov chain when the loss rate is small. Ott, Kemperman and Mathis [20] give an analysis of the evolution of the congestion window size via a differential equation perturbed by a Poisson process. Dumas, Guillemin and Robert [8] provide rigorous convergence results and explicit expressions of the stationary distributions for the congestion avoidance regime when packet losses are independent and the loss rate tends to 0.

In this paper the behavior of a persistent TCP connection experiencing packet losses is investigated. Instead of assuming that packet losses are independent from one packet to another, the case when packet losses occur in clumps is investigated. This model can be explained by the fact that a loss of a packet is due to the overflow of some buffer somewhere in the network. In this situation, very likely, losses will continue to occur for some time (until the TCP connections involved decrease their transmission rate). This model for packet losses has been validated by recent measurements made by Paxson [22] on the loss process affecting TCP connections in the Internet. (See also [6] and [25].) Some papers considered analytical models describing the case of bursts of losses for TCP connections. Misra, Gong and Towsley [18] analyzed, in a setting similar to [20], a representation of the sequence of the congestion window sizes as an $M/G/1$ queue. This $M/G/1$ representation is also used in [2] to study grouped packet losses. In these papers the probability of a loss of a packet in a congestion window of size W is independent of W; this is not the case for the model considered here. The more it is sent into the network, the more likely a packet loss occurs.

On the probabilistic side, it is shown that the so-called *exponential functionals* of Lévy processes describe the asymptotic behavior of AIMD algorithms. Exponential functionals have received much attention recently, motivated by applications in mathematical finance (the Lévy process is a Brownian motion with drift in this setting) or in statistical physics. See Yor's book [28] on this subject. In the case of TCP, the corresponding Lévy processes are compound Poisson processes. The calculation of the density function of these random variables turns out to be quite intricate. Analytically, the natural framework is the q-calculus (the Appendix gives a brief presentation of this topic). In this setting, the distributions of some of these exponential functionals are related to q-hypergeometric functions.

The paper studies the asymptotic behavior of the TCP connection when the loss rate converges to 0. It is organized as follows.

In Section 2 the basic convergence results are established, they are straightfor-
ward generalizations of analogous results proved in the case of independent losses
in [8]. The main results of the paper, in Sections 3 and 4, concern the asymptotic
distributions which are more delicate to investigate than for the independent loss
model. In Section 3 the asymptotic invariant distribution of the congestion window
size is analyzed and the exponential functionals are introduced. The density func-
tion of these random variables is expressed in terms of a functional of a random
walk and, in some cases, as q-hypergeometric function. An explicit expression for
their fractional moments is given. In Section 4 the asymptotic throughput is in-
vestigated. The fractional moment of order $1/2$ of the corresponding exponential
functional obtained in Section 3 is used to get an explicit expression of the asymp-
totic throughput. The rest of the section is devoted to the impact of correlations on
the throughput. It is shown that, for concave ordering, the throughput is a nonin-
creasing function of the distribution of the number of losses in a clump (a group of
local losses). In particular the model of independent packet losses turns out to be a
pessimistic model since it underestimates the real performances of TCP. Section 5
concludes the paper by a discussion of some additional features of TCP which are
not represented in the stochastic model.

2. A model with correlated losses. It is assumed that a data connection
transmits packets into a network by means of an AIMD algorithm with additive
increase factor 1 and multiplicative decrease equal to $\delta < 1$. Let W_n^α denote the
congestion window size over the nth RTT (Round Trip Time) interval, that is,
the total number of packets sent during this time interval. The evolution of the
process (W_n^α) is given by

$$(2.1) \quad W_{n+1}^\alpha = \begin{cases} W_n^\alpha + 1, & \text{when none of the } W_n^\alpha \text{ packets is lost,} \\ \max(\lfloor \delta W_n^\alpha \rfloor, 1), & \text{otherwise.} \end{cases}$$

To complete the presentation of the model, the loss process of the packets has
to be described. In the noncorrelated case considered in [8], each packet has a
probability $1 - \exp(-\alpha)$ of being lost. For $n \geq 1$, if t_n^α denotes the index of nth
which is lost, the independence assumption implies that the sequence $(t_{n+1}^\alpha - t_n^\alpha)$
is i.i.d. having a geometric distribution with parameter $\exp(-\alpha)$. Hence, when
α is small, then $t_{n+1}^\alpha - t_n^\alpha \sim E_n/\alpha$, where E_n is exponentially distributed with
parameter 1. Asymptotically, the loss process can thus be described as a Poisson
process.

This noncorrelated loss process is not completely realistic since it does not
take into account the fact that a loss is due to an overflow of some buffer in
the network. Therefore, after a loss, subsequent losses are more likely. On the
other hand, since the state of the network changes quite rapidly, the network
"forgets" the past quickly; the i.i.d. assumption for $(t_{n+1}^\alpha - t_n^\alpha)$ can be plausible
provided that these quantities are not too small. Paxson [22] and Zhang, Paxson
and Shenker [29] showed through measurements that the loss process can in fact
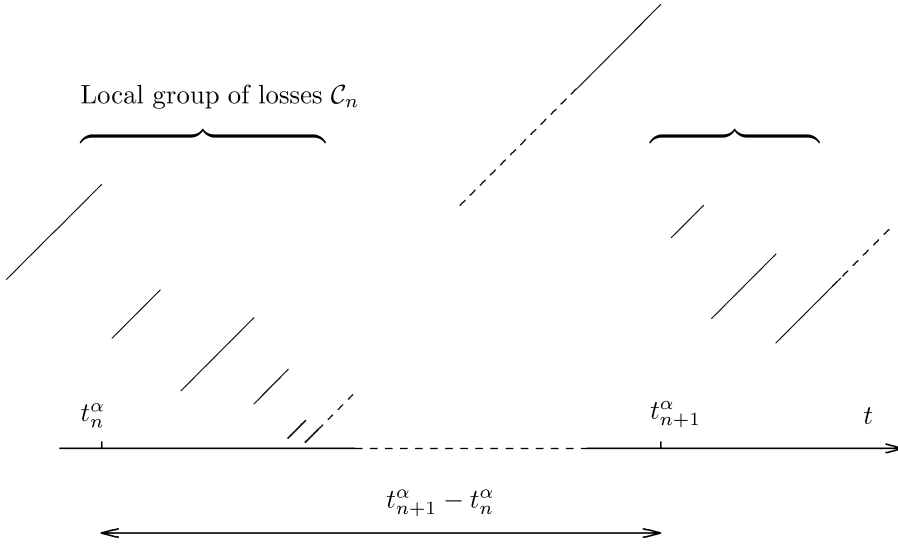
FIG. 1. *Evolution of the congestion window size between two groups of losses.*

be described as follows: when a packet loss occurs, several packets are also lost during the following RTT intervals. After this clump of losses, the next packet loss will occur as in the independent model. Mathematically, this can be described as follows: the indexes of the packets involved in the nth clump are given by the set $t_n^\alpha + \mathcal{C}_n$ (see Figure 1), where:

(a) (t_n^α) is the sequence considered in the noncorrelated case.
(b) The set $\mathcal{C}_n$ is a finite subset of $\mathbb{N}$ containing 0 and the sequence $(\mathcal{C}_n)$ is i.i.d.

Therefore, the quantity t_n^α is the index of the first packet lost in the nth group of losses. If $x \in \mathcal{C}_n$, the packet with index $(t_n^\alpha + x)$ belongs to the nth group of losses. It is assumed that the distribution of the $\mathcal{C}_n$'s does not depend on α. When α is small, the distance between two group of losses is large since $\mathbb{E}(t_{n+1}^\alpha - t_n^\alpha) = O(1/\alpha)$. This represents the fact, observed by Paxson, that multiple losses may occur locally. The i.i.d. assumption of the sequence $(\mathcal{C}_n)$ is a consequence of the rapid changes of the network between two groups of losses. In particular, if the cardinality of $\mathcal{C}_n$ is denoted by X_n, the sequence (X_n) is i.i.d. As we shall see later this sequence (X_n) gives a measure of the correlation of packet losses. The packet loss rate of this stochastic model is thus equivalent to $\alpha\mathbb{E}(X_1)$ as α gets small.

Asymptotically, at the packet level, the loss process can thus be described as a Poisson process with clumps, that is, a standard Poisson process with "clouds" around each of its points. This representation of the occurrences of rare events is quite universal in probability theory. Aldous' book [1] illustrates, through a large collection of examples, the generality of this description.

REMARKS.

1. The case $\mathcal{C}_1 \equiv \{0\}$ corresponds to the uncorrelated case. One packet is lost in each group of losses.
2. The i.i.d. assumption $(\mathcal{C}_n)$ is a consequence of the fact that the network forgets; at "time" t_{n+1}^α the events $t_n^\alpha + \mathcal{C}_n$ have been forgotten, in particular, $\mathcal{C}_{n+1}$ is independent of $\mathcal{C}_n$.
3. The set $\mathcal{C}_n$ could depend on α, provided that the location of its last element is negligible compared to $1/\alpha$. (Recall that the set $t_n^\alpha + \mathcal{C}_n$ has to be far way from t_{n+1}^α.) For simplicity the independence with respect to α is assumed, it is easily checked that this is not really restrictive.

With such a loss process the sequence (W_n^α) does not necessarily have the Markov property. The i.i.d. property of the sequence $(\mathcal{C}_n)$ shows nevertheless that, if l_n is the largest element of $\mathcal{C}_n$, the embedded chain $(V_n^\alpha) = (W_{t_n^\alpha + l_n}^\alpha)$ (the sequence of congestion window sizes at the end of groups of losses) is still Markov. The next proposition shows that, properly renormalized, the transitions of this Markov chain converge.

For $m \geq 1$, the variable G_m^α is defined as

$$G_m^\alpha = \inf\{k : W_{k+1} = \max(\lfloor \delta W_k^\alpha \rfloor, 1)\},$$

with $W_0 = m$. When the initial window size is m, G_m^α is the number of successful RTT intervals until the next packet loss. Thus, for $n \geq 0$, if $W_{t_n^\alpha + l_n} = m$, then

$$W_{t_{n+1}^\alpha} \overset{\text{dist.}}{=} \max(\lfloor \delta(m + G_m^\alpha) \rfloor, 1).$$

PROPOSITION 1. *For $x > 0$, as α goes to 0, the random variable $\sqrt{\alpha}\, G_{\lfloor x/\sqrt{\alpha} \rfloor}^\alpha$ converges in distribution to a nonnegative random variable $\overline{G}_x$ such that for $y \geq 0$,*

$$(2.2) \qquad\qquad \mathbb{P}(\overline{G}_x \geq y) = \exp(-xy - y^2/2).$$

If $V_0^\alpha = \lfloor x/\sqrt{\alpha} \rfloor$ then, as α tends to 0, the random variable $\sqrt{\alpha}\, V_1^\alpha$ converges in distribution to $\overline{V}_1$ with

$$(2.3) \qquad\qquad \overline{V}_1 = \delta^{X_1}(x + \overline{G}_x),$$

where X_1 and $\overline{G}_x$ are independent random variables.

PROOF. The first part of the proposition is easily seen; it has been proved in [8]. If $V_0^\alpha = \lfloor x/\sqrt{\alpha} \rfloor$ at $t_0^\alpha + l_0$, where l_0 is the last element of $\mathcal{C}_0$, then at t_1^α, one has $W_{t_1^\alpha}^\alpha = \lfloor \delta(V_0^\alpha + G_{V_0^\alpha}^\alpha) \rfloor$. Hence $\sqrt{\alpha}\, W_{t_1^\alpha}^\alpha$ converges in distribution to $\delta(x + \overline{G}_x)$ as α tends to 0.

The factor δ^{X_1} in (2.3) is a consequence of the X_1 losses occurring in the clump of losses $t_1^\alpha + \mathcal{C}_1$, with the underlying property that the window size does not grow significantly (with respect to $1/\alpha$) during that period.

More rigorously, if r_1 is the second point of $t_1^\alpha + \mathcal{C}_1$, then

$$W_{r_1}^\alpha = \lfloor \delta(W_{t_1^\alpha}^\alpha + r_1 - t_1^\alpha)\rfloor.$$

Since $r_1 - t_1^\alpha \le l_1$ and l_1 does not depend on α, the following convergence in distribution holds

$$\lim_{\alpha \to 0} \sqrt{\alpha} W_{r_1}^\alpha = \delta \lim_{\alpha \to 0} \sqrt{\alpha} W_{t_1^\alpha}^\alpha \overset{\text{dist.}}{=} \delta^2(x + \overline{G}_x).$$

By induction on the number of points of $t_1^\alpha + \mathcal{C}_1$, one finally gets

$$\lim_{\alpha \to 0} \sqrt{\alpha} V_1^\alpha \overset{\text{dist.}}{=} \delta^{X_1}(x + \overline{G}_x).$$

The proposition is proved. $\square$

The model considered here does not distinguish between the different kinds of losses: losses due to a timeout or losses detected by reception a triple dupliquate acknowledgment. (See [23].) In the present implementations of TCP, when a timeout occurs, the congestion window size is set to 1 and the slow start procedure is used instead of the AIMD algorithm. For the moment this part of TCP is not considered, we shall see that it can be included without any problem (see Section 5.2).

The above proposition shows that, if $V_0^\alpha = \lfloor x/\sqrt{\alpha}\rfloor$, the Markov chain $(\sqrt{\alpha} V_n^\alpha)$ converges to the continuous state space Markov chain $(\overline{V}_n)$ with $\overline{V}_0 = x$ and

(2.4) $$\overline{V}_{n+1} = \delta^{X_n}(\overline{V}_n + \overline{G}_{\overline{V}_n})$$

for $n \in \mathbb{N}$.

This result established, the convergence results are stated without proof. Proofs are exactly the same as in the uncorrelated case investigated in [8]. The major difference, this is the main point of the paper, is the fact that closed form expressions of the limiting distributions are much more difficult to derive as it will be seen in the following sections.

THEOREM 2. *When α tends to 0, the invariant distribution of the Markov chain $(\sqrt{\alpha} V_n^\alpha)$ converges in distribution to the invariant distribution $\overline{V}_\infty$ of the Markov chain $(\overline{V}_n)$.*

With a slight abuse of notation, the expression "the invariant distribution $\overline{V}_\infty$" means "a random variable $\overline{V}_\infty$ whose distribution is invariant for the Markov chain."

PROPOSITION 3. *The invariant distribution $\overline{V}_\infty$ of the continuous state space Markov chain $(\overline{V}_n)$ satisfies the following identities:*

(2.5) $$\overline{V}_\infty^2 \overset{\text{dist.}}{=} \delta^{2X_1}(\overline{V}_\infty^2 + 2E_1),$$

where X_1, E_1 and $\overline{V}_\infty$ are independent random variables, E_1 being exponentially distributed with parameter 1.

This proposition is a simple consequence of the elementary identity in distribution, for $x \geq 0$,

$$(2.6) \qquad (x + \overline{G}_x)^2 \stackrel{\text{dist.}}{=} x^2 + 2E_1.$$

(See [8].)

Similarly, convergence results can also be obtained for the original sequence (W_n^{α}).

PROPOSITION 4.   *If* $\lim_{\alpha \to 0} \sqrt{\alpha} W_0^{\alpha} = \overline{w}$, *then*

$$(W^{\alpha}(t)) = \left(\sqrt{\alpha} W_{\lfloor t / \sqrt{\alpha} \rfloor}^{\alpha} \right)$$

converges in distribution to the Markov process $(\overline{W}(t))$ *such that* $\overline{W}(0) = \overline{w}$ *and with the infinitesimal generator given by*

$$(2.7) \qquad \Omega(f)(x) = f'(x) + x \int_{\mathbb{R}_+} \left(f(\delta^u x) - f(x) \right) X_1(du)$$

for any C^1-*function* f *on* $\mathbb{R}_+$, *where* $X_1(dx)$ *denotes the distribution of* X_1 *on* $\mathbb{N}$.

3. The exponential functional of a compound Poisson process.
In this section the distribution of the random variable I, solution to the equation

$$(3.1) \qquad I \stackrel{\text{dist.}}{=} \beta^{X_1} I + E_0,$$

is investigated, where $\beta \in [0, 1[$, the variables E_0, I and X_1 are independent, E_0 is an exponentially distributed random variable with parameter 1 and X_1 is some non-negative random variable (not necessarily integer valued) such that $\mathbb{P}(X_1 > 0) = 1$.

In view of (2.5), if β is δ^2, it is easily seen that I and $\overline{V}_{\infty}^2 / 2 + E_0$ have the same distribution. By iterating (3.1), the variable I can be represented as

$$(3.2) \qquad I = \sum_{n=0}^{+\infty} \beta^{S_n} E_n,$$

where (E_n) is an i.i.d. sequence of exponentially distributed random variables with parameter 1 and $(S_n) = (X_1 + \cdots + X_n)$ is the random walk associated to the i.i.d. sequence (X_n). If $(N(t))$ is a Poisson process with parameter 1 such that, for $n \geq 0$, the distance between the $(n+1)$st point and the nth point is E_{n+1} and if $(\xi(t))$ is the compound Poisson process

$$\xi(t) = \log(1/\beta) \sum_{k=1}^{N(t)} X_k,$$

it is easily seen that (3.2) of I can be written as

$$(3.3) \qquad I = \int_0^{+\infty} e^{-\xi(t)} dt.$$

The variable I is the *exponential functional* associated to the Lévy process $(\xi(t))$. It occurs naturally in mathematical finance (Asian options) and in many other fields; the Lévy process ξ is generally a Brownian motion with drift. In this setting, when the variable I is introduced by (3.3), Carmona, Petit and Yor [7] proved, that the density of I is the solution to an integro-differential equation. In other words, they showed that the distribution of the random variable I is the invariant distribution of some Markov process. Notice we followed the reverse path in our analysis. Finally, let us mention that a lot of work has been done when the Lévy process is related to a continuous diffusion. Yor [28] surveys these questions; see also [26], Chapter 8 and [27], Section 15.4 for a more theoretical point of view.

PROPOSITION 5. *For $\lambda \geq 0$, the Laplace transform of the variable I defined by (3.2) is given by*

$$(3.4) \qquad \mathbb{E}(e^{-\lambda I}) = \mathbb{E}\left(\prod_{n=0}^{+\infty} \frac{1}{1 + \lambda \beta^{S_n}} \right),$$

where $(S_n) = (X_1 + \cdots + X_n)$ is the random walk associated to the i.i.d. sequence (X_n).

If there exists some $\varepsilon > 0$ such that $\mathbb{P}(X_1 \geq \varepsilon) = 1$, the density h of I is given by, for $x \geq 0$,

$$(3.5) \qquad h(x) = C \sum_{n=0}^{+\infty} \mathbb{E}\left(\prod_{k=1}^{n} \frac{1}{1 - \beta^{-S_k}} \beta^{-S_n} e^{-\beta^{-S_n} x} \right),$$

with $C = \mathbb{E}(1/\prod_{n=1}^{+\infty}(1 - \beta^{S_n}))$.

PROOF. Representation (3.4) of the Laplace transform of I is obtained directly from (3.2). The random variable

$$H(\lambda) = \prod_{n=0}^{+\infty} \frac{1}{1 + \lambda \beta^{S_n}}$$

is a (random) meromorphic function of λ on $\mathbb{C}$. The assumption $\mathbb{P}(X_1 \geq \varepsilon) = 1$ implies that the function H has only simple poles located in $\{-\beta^{-S_n} : n \geq 0\}$. For $n \geq 0$, its residue at $-\beta^{-S_n}$ is given by

$$\prod_{k=0}^{n-1} \frac{1}{1 - \beta^{S_k - S_n}} \beta^{-S_n} \prod_{k=n+1}^{\infty} \frac{1}{1 - \beta^{S_k - S_n}},$$

therefore $H(\lambda) = \sum_{n \geq 0} R_n(\lambda)$, with

$$R_n(\lambda) = \frac{1}{1 + \lambda \beta^{S_n}} \prod_{k=0}^{n-1} \frac{1}{1 - \beta^{S_k - S_n}} \prod_{k=n+1}^{\infty} \frac{1}{1 - \beta^{S_k - S_n}}.$$

The i.i.d. property of the sequence (X_n) shows that

$$\mathbb{E}(R_n(\lambda)) = \mathbb{E}\left(\frac{1}{1+\lambda\beta^{S_n}} \prod_{k=0}^{n-1} \frac{1}{1-\beta^{S_k-S_n}}\right) \mathbb{E}\left(\prod_{k=n+1}^{\infty} \frac{1}{1-\beta^{S_k-S_n}}\right)$$

(3.6)

$$= \mathbb{E}\left(\frac{1}{1+\lambda\beta^{S_n}} \prod_{k=1}^{n} \frac{1}{1-\beta^{-S_k}}\right) \mathbb{E}\left(\prod_{k=1}^{\infty} \frac{1}{1-\beta^{S_k}}\right).$$

Due to the assumption on the distribution of X_1, for $\lambda \geq 0$, one gets the inequality

$$|R_n(\lambda)| \leq \prod_{k=1}^{n-1} \frac{\beta^{S_n-S_k}}{1-\beta^{S_n-S_k}} \prod_{k=n+1}^{\infty} \frac{1}{1-\beta^{S_k-S_n}} \leq \prod_{k=1}^{n-1} \frac{\beta^{k\varepsilon}}{1-\beta^{k\varepsilon}} \prod_{k=1}^{\infty} \frac{1}{1-\beta^{k\varepsilon}}.$$

Fubini's theorem therefore shows that

$$\mathbb{E}(\exp(-\lambda I)) = \mathbb{E}(H(\lambda)) = \sum_{n\geq 0} \mathbb{E}(R_n(\lambda)).$$

Identity (3.6) gives that $\mathbb{E}(R_n(\lambda))$ is the Laplace transform of the density $r_n(x)$, with

$$r_n(x) = \mathbb{E}\left(\beta^{-S_n} e^{-\beta^{-S_n} x} \prod_{k=1}^{n} \frac{1}{1-\beta^{-S_k}}\right) \mathbb{E}\left(\prod_{k=1}^{\infty} \frac{1}{1-\beta^{S_k}}\right),$$

the density of I can be thus expressed as the sum of the r_n's. The proposition is proved. $\square$

Representation (3.5) can be used to obtain explicit expressions for the density of I only when the distributions of some functionals of the random walk (S_n) are known. In general, this is not the case (see the examples below). Formula (3.5) is nevertheless useful to get numerical expressions since the general term of the series converges rapidly.

EXAMPLES.

1. *The case $X_1 \equiv 1$.* This is the situation considered in [7] and [8] (with $\beta = \delta^2$). Since in this case $S_n = n$ for all $n \geq 0$, (3.5) gives

(3.7) $$h(x) = \frac{1}{\prod_{n=1}^{+\infty}(1-\beta^n)} \sum_{n=0}^{+\infty} \frac{1}{\prod_{k=1}^{n}(1-\beta^{-k})} \beta^{-n} e^{-\beta^{-n}x}.$$

2. *The distribution of X_1 is exponential with parameter μ.* According to (3.4), the Laplace transform of I at $\lambda \geq 0$ is given by

$$\mathbb{E}(e^{-\lambda I}) = \frac{1}{1+\lambda} \mathbb{E}\left(\exp\left(-\sum_{n=1}^{+\infty} \log(1+\lambda\beta^{S_n})\right)\right).$$

Clearly enough (S_n) is a Poisson point process on $\mathbb{R}^+$ with parameter μ, using the expression of the Laplace transform of a Poisson point process (see [19] or [16] for example). The expected value in the right-hand side of the last equation is thus given by

$$\exp\left(-\int_0^{+\infty}\left(1-\frac{1}{1+\lambda\beta^x}\right)\mu\,dx\right)$$

$$=\exp\left(\mu/\log(\beta)\int_0^1\frac{\lambda}{1+\lambda u}\,du\right)=\left(\frac{1}{1+\lambda}\right)^{-\mu/\log(\beta)},$$

hence,

$$\mathbb{E}(e^{-\lambda I})=\left(\frac{1}{1+\lambda}\right)^{1-\mu/\log(\beta)}.$$

The density of I is therefore the Gamma density function with parameter $(1-\mu/\log(\beta))$,

$$h(x)=\frac{x^{-\mu/\log(\beta)}}{\Gamma(1-\mu/\log(\beta))}e^{-x},\qquad x\geq 0.$$

See [14].

3.1. *The fractional moments.*   In general, a useful explicit expression of the distribution of I is not easy to derive. It turns out that the moments of I can be expressed quite easily, including the fractional moments of I. This is clearly useful since the stationary window size $\overline{V}_\infty$ can be expressed with the square root of the random variable I; indeed, (2.5) shows that

$$\overline{V}_\infty^2/2 \stackrel{\text{dist.}}{=} \delta^{2X_1}\left(\overline{V}_\infty^2/2+E_1\right)\stackrel{\text{dist.}}{=}\delta^{2X_1}I.$$

Section 4 uses a fractional moment of I to derive an explicit expression of the throughput of the AIMD algorithm.

PROPOSITION 6 (A recursive formula for the moments of I).   *For any $s \in \mathbb{R}$,*

(3.8)
$$\mathbb{E}(I^{s-1})=\frac{1-\mathbb{E}(\beta^{sX_1})}{s}\mathbb{E}(I^s),$$

the moment of order s of I is finite if $\mathbb{E}(\beta^{(s+1)X_1}) < +\infty$.

For $s \geq 0$, relationship (3.8) is due to Carmona, Petit and Yor [7]. Notice that the condition

$$\mathbb{E}(\beta^{(s+1)X_1}) < +\infty$$

is dummy for $s > -1$ since $\beta < 1$.

PROOF OF PROPOSITION 6.   For $\lambda \geq 0$, denote by $\psi(\lambda)$ the Laplace transform of I at λ, (3.1) gives

$$(3.9) \qquad \psi(\lambda) = \frac{1}{1+\lambda} \mathbb{E}\big(\psi(\lambda \beta^{X_1})\big),$$

the Mellin transform of ψ is

$$\psi^*(s) = \int_0^{+\infty} \psi(\lambda) \lambda^{s-1} \, d\lambda.$$

Since ψ is bounded, ψ^* is defined for $\Re(s) > 0$. (See [11] for a survey on Mellin transform methods.) By using the definition of ψ,

$$\psi^*(s) = \mathbb{E}\left(\int_0^{+\infty} e^{-\lambda I} \lambda^{s-1} \, d\lambda \right)$$

$$(3.10)$$

$$= \mathbb{E}\left(\frac{1}{I^s} \right) \int_0^{+\infty} e^{-\lambda} \lambda^{s-1} \, d\lambda = \mathbb{E}\left(\frac{1}{I^s} \right) \Gamma(s).$$

On the other hand, (3.9),

$$(1+\lambda)\psi(\lambda) = \mathbb{E}\big(\psi(\lambda \beta^{X_1})\big),$$

becomes, via Mellin transform,

$$\psi^*(s) + \psi^*(s+1) = \mathbb{E}(\beta^{-sX_1})\psi^*(s),$$

and, by (3.10),

$$(3.11) \qquad \mathbb{E}\left(\frac{1}{I^{s+1}} \right) = \frac{\mathbb{E}(\beta^{-sX_1}) - 1}{s} \mathbb{E}\left(\frac{1}{I^s} \right).$$

This relationship extends on $\mathbb{R}_-$ formula (3.8) obtained by Carmona, Petit and Yor [7] for $s \in \mathbb{R}_+$. If I has a finite, moment of order $-s$ and $\mathbb{E}(\beta^{-sX_1})$ is finite, then I has a finite moment of order $-s-1$. Since all the positive moments of I are finite (see [7] for example), by induction one gets that $\mathbb{E}(1/I^s)$ is finite when $\mathbb{E}(\beta^{-(s-1)X_1})$ is finite. The proposition is then proved.   □

PROPOSITION 7.   *For any $s \in \mathbb{R}$, $-s \notin \mathbb{N} \setminus \{0\}$, if $\mathbb{E}(\beta^{(s+1)X_1}) < +\infty$ and*

$$\mathbb{E}\left(\frac{1}{1 - \beta^{X_1}} \right) < +\infty,$$

then the moment of order s of the variable I can be expressed as

$$(3.12) \qquad \mathbb{E}(I^s) = \Gamma(s+1) \prod_{k=1}^{+\infty} \frac{\phi(s+k)}{\phi(k)},$$

where $\phi(u) = 1 - \mathbb{E}(\beta^{uX_1})$ for $u \geq \min(s, 0)$.

When $s \in \mathbb{N}$, (3.12) gives that

$$(3.13) \qquad \mathbb{E}(I^s) = \frac{s!}{\prod_{k=1}^s (1 - \mathbb{E}(\beta^{kX_1}))}$$

and when $-s \in \mathbb{N} \setminus \{0\}$, (3.12) can be continued by using the fact that the Gamma function has a simple pole at $s + 1$ whose residue is $(-1)^{-s}/(-s)!$ so that

$$(3.14) \qquad \mathbb{E}\left(\frac{1}{I^s}\right) = \frac{\prod_{k=s}^{-1}(\mathbb{E}(\beta^{kX_1}) - 1)}{s!} \mathbb{E}(X_1).$$

Relationships (3.13) and (3.14) have already been remarked in [7] and [5]. Identity (3.12) has been obtained independently by Bertoin, Biane and Yor [4] when $X_1 \equiv 1$.

For a general Lévy process $(\xi(t))$, ϕ is the Lévy–Khintchine exponent defined by

$$\mathbb{E}\big(e^{-s\xi(1)}\big) = e^{-\phi(s)},$$

for $s \geq 0$. Under mild assumptions on ϕ (see the proof of Proposition 7), (3.12) should hold for the corresponding exponential functional I.

PROOF OF PROPOSITION 7. First note that, for N sufficiently large

$$\sum_{i=N}^{+\infty} |\log(\phi(k))| = \sum_{i=N}^{+\infty} \big|\log\big(1 - \mathbb{E}(\beta^{kX_1})\big)\big|$$

$$\leq 2 \sum_{i=1}^{+\infty} \mathbb{E}(\beta^{kX_1}) \leq 2\mathbb{E}\left(\frac{1}{1 - \beta^{X_1}}\right) < +\infty,$$

hence, the right-hand side of (3.12) is well defined. Denote by ψ the function

$$\psi(s) = \mathbb{E}(I^{s-1}) \prod_{k=1}^{+\infty} \frac{\phi(k)}{\phi(s+k-1)},$$

according to (3.8) the function ψ satisfies the functional equation

$$\psi(s+1) = s\psi(s),$$

for any $s > 0$. This relationship is also satisfied by the classical Gamma function. Since $\psi(1) = 1$, to prove that ψ is indeed Γ, Bohr–Mollerup's theorem (see [3]) shows that it is sufficient to prove that ψ is log-convex, that is, that $\log(\psi)$ is a convex function on $\mathbb{R}_+^* = \mathbb{R}_+ \setminus \{0\}$. For $s > 0$,

$$\log(\psi(s)) = \log\left(\prod_{k=1}^{+\infty} \phi(k)\right) + \log \mathbb{E}(I^{s-1}) + \sum_{k=1}^{+\infty} -\log\big(1 - \mathbb{E}\big(\beta^{(s+k-1)X_1}\big)\big),$$

since I^u is integrable for any $u > -1$, it is easily seen that for $s > 0$, the variables $I^{s-1} \log I$ and $I^{s-1}(\log I)^2$ are integrable. The function $s \to \log \mathbb{E}(I^{s-1})$ is thus twice differentiable and its second derivative is given by

$$\frac{\mathbb{E}(I^{s-1}(\log I)^2)}{\mathbb{E}(I^{s-1})} - \left(\frac{\mathbb{E}(I^{s-1} \log I)}{\mathbb{E}(I^{s-1})} \right)^2;$$

it is nonnegative by Cauchy–Schwarz's inequality. Similarly, for $k \geq 1$ the function

$$s \to -\log(1 - \mathbb{E}(\beta^{(s+k-1)X_1}))$$

is also convex on $\mathbb{R}_+^*$. The function ψ is therefore log-convex, hence $\psi = \Gamma$ on $\mathbb{R}_+^*$. Relationship (3.12) holds on $\mathbb{R}_+^*$ and it is partially extended on $\mathbb{R}_-$ by using (3.8). The proposition is proved. $\square$

3.2. *The density function as a q-hypergeometric function.*   The integer moments of the variable I can be naturally used to get a representation of the Laplace transform of I.

PROPOSITION 8.   *The Laplace transform of the random variable I is given by, for $\lambda \in [0, 1)$,*

$$(3.15) \qquad \mathbb{E}(e^{-\lambda I}) = \sum_{n=0}^{+\infty} \frac{(-\lambda)^n}{\prod_{k=1}^n (1 - \mathbb{E}(\beta^{kX_1}))}.$$

PROOF.   See also [5]. The representation

$$\mathbb{E}(e^{-\lambda I}) = \sum_{n=0}^{+\infty} \mathbb{E}(I^n) \frac{(-\lambda)^n}{n!}$$

is valid when the above series converge and Carleman's criterion is verified (see [9]), that is, if

$$\sum_{n=0}^{+\infty} \mathbb{E}(I^{2n})^{-1/2n} = +\infty,$$

which is easily checked by using (3.13). $\square$

If the random variable X_1 has a rational generating function, that is, there exist two polynomials P and Q such that $\mathbb{E}(z^{X_1}) = P(z)/Q(z)$, then for some $a_1, \ldots, a_M, b_1, \ldots, b_N \in \mathbb{C}$,

$$1 - \mathbb{E}(z^{X_1}) = \frac{(1 - z) \prod_{j=1}^N (1 - b_j z)}{\prod_{i=1}^M (1 - a_i z)}.$$

A direct consequence of (3.15) is the following representation for the Laplace transform of random variable I:

$$(3.16) \qquad \mathbb{E}(e^{-\lambda I}) = \sum_{n=0}^{+\infty} \frac{(a_1\beta; \beta)_n \cdots (a_M\beta; \beta)_n}{(b_1\beta; \beta)_n \cdots (b_N\beta; \beta)_n} \frac{(-\lambda)^n}{(\beta; \beta)_n},$$

where, for $x \in \mathbb{C}$, $q \in [0, 1[$, $(x; q)_n$ is defined by

$$(x; q)_n = (1 - x)(1 - xq)(1 - xq^2) \cdots (1 - xq^{n-1})$$

for $k \geq 1$ and $(a; q)_0 = 1$. Expression (3.16) for the Laplace transform can be transformed so that it can be expressed as a q-hypergeometric functions. See definition (A.1) and some basic identities in the Appendix. This suggests that q-calculus is the natural setting to study the density of exponential functionals for discontinuous Lévy processes. See [4] for some developments in this setting when the process is purely Poisson. Different cases are now analyzed.

(1) *The shifted geometric distribution.* First, let us consider the case when X_1 has a shifted geometric distribution, that is, for $a < 1$ and $n \geq 1$,

$$\mathbb{P}(X_1 = n) = a^{n-1}(1 - a).$$

If a is not a power of β, that is, $a \notin \{\beta^p : p \geq 1\}$. For $|z| \leq 1$,

$$1 - \mathbb{E}(z^X) = \frac{1 - z}{1 - az},$$

and from (3.16) one gets the relationship

$$\mathbb{E}(e^{-\lambda I}) = \sum_{n=0}^{+\infty} \frac{(a\beta; \beta)_n}{(\beta; \beta)_n}(-\lambda)^n.$$

The q-Binomial theorem (Theorem 18 recalled in the Appendix) gives

$$\mathbb{E}(e^{-\lambda I}) = \frac{(-\lambda a\beta; \beta)_\infty}{(-\lambda; \beta)_\infty},$$

therefore, Laplace transform I has simple poles at points $-\beta^{-n}$, $n \geq 0$, and the residue at point $-\beta^{-n}$ is

$$\frac{\beta^{-n}(a\beta^{-n+1}; \beta)_\infty}{(\beta; \beta)_\infty \prod_{k=1}^n (1 - \beta^{-k})} = \frac{\beta^{-n}(a\beta^{-n+1}; \beta)_\infty}{(\beta; \beta)_\infty (1/\beta; 1/\beta)_n}.$$

The density h of the distribution of the random variable I is thus given by, for $x \geq 0$,

$$(3.17) \qquad h(x) = \frac{1}{(\beta; \beta)_\infty} \sum_{n=0}^{+\infty} \frac{(a\beta^{-n+1}; \beta)_\infty}{(1/\beta; 1/\beta)_n} \beta^{-n} e^{-\beta^{-n} x}.$$

For $a = 0$, $X = 1$ a.s. and the two probability distributions defined by (3.17) and (3.7) coincide.

Second, if $a = \beta^p$ for some $p \geq 1$, the Laplace transform of I is a rational function given by

$$\mathbb{E}(e^{-\lambda I}) = \frac{1}{(-\lambda; \beta)_p},$$

thus

$$h(x) = \sum_{n=0}^{p} \frac{1}{(1/\beta; 1/\beta)_n (\beta, \beta)_{p-n}} \beta^{-n} e^{-\beta^{-n} x}.$$

(2) *A two-valued random variable X_1.* Here, $\mathbb{P}(X_1 = 1) = p$ and $\mathbb{P}(X_1 = 2) = 1 - p$, then $1 - \mathbb{E}(z^{X_1}) = (1 - z)(1 + (1 - p)z)$. The Laplace transform of the random variable I is thus given by

$$\mathbb{E}(e^{-\lambda I}) = \sum_{n=0}^{+\infty} \frac{(-\lambda)^n}{(\beta; \beta)_n (-\beta(1 - p); \beta)_n}.$$

The elementary relationship

$$(z; q)_n = \frac{(z; q)_\infty}{(q^n z; q)_\infty}$$

gives

$$\mathbb{E}(e^{-\lambda I}) = \frac{1}{(-(1 - p)\beta; \beta)_\infty} \sum_{n=0}^{+\infty} \frac{(-(1 - p)\beta^{n+1}; \beta)_\infty}{(\beta; \beta)_n} (-\lambda)^n.$$

From the first Euler's identity (A.3) in the Appendix, one gets

$$\mathbb{E}(e^{-\lambda I}) = \frac{1}{(-(1 - p)\beta; \beta)_\infty} \sum_{n=0}^{+\infty} \sum_{m=0}^{+\infty} \beta^{m(m-1)/2} \frac{(1 - p)^m \beta^{(n+1)m}}{(\beta; \beta)_m} \frac{(-\lambda)^n}{(\beta; \beta)_n}$$

$$= \frac{1}{(-(1 - p)\beta; \beta)_\infty} \sum_{m=0}^{+\infty} \beta^{m(m-1)/2} \frac{(1 - p)^m \beta^m}{(\beta; \beta)_m} \sum_{n=0}^{+\infty} \frac{(-\lambda \beta^m)^n}{(\beta; \beta)_n},$$

and then, using the second Euler identity (A.4),

$$\mathbb{E}(e^{-\lambda I}) = \frac{1}{(-(1 - p)\beta; \beta)_\infty} \sum_{m=0}^{+\infty} \beta^{m(m-1)/2} \frac{(1 - p)^m \beta^m}{(\beta; \beta)_m (-\lambda \beta^m; \beta)_\infty}$$

$$= \frac{1}{(-(1 - p)\beta; \beta)_\infty (-\lambda; \beta)_\infty} \sum_{m=0}^{+\infty} \beta^{m(m+1)/2} \frac{(-\lambda; \beta)_m}{(\beta; \beta)_m} (1 - p)^m.$$

It follows that the Laplace transform of the random variable I has simple poles at the points $-\beta^{-n}$, $n \geq 0$. The residue of the Laplace transform at $-\beta^{-n}$ is equal to $r_n \beta^{-n}$, with

$$r_n = \frac{1}{(-(1-p)\beta; \beta)_\infty (\beta; \beta)_\infty} \sum_{m=0}^{n} \frac{(-1)^m (1-p)^m}{(1/\beta; 1/\beta)_m (1/\beta; 1/\beta)_{n-m}}.$$

The density h of the random variable I is thus given by, for $x \geq 0$,

$$(3.18) \qquad h(x) = \sum_{n=0}^{+\infty} r_n \beta^{-n} e^{-\beta^{-n} x}.$$

(3) *A random variable X_1 with a rational generating function.* The above examples can be generalized in the following manner.

PROPOSITION 9. *If for $a \in \mathbb{C}$ and $b_1, \ldots, b_N \in \mathbb{C}$ such that, for $|z| \leq 1$, the generating function of X_1 is given by*

$$1 - \mathbb{E}(z^{X_1}) = \frac{(1-z) \prod_{i=1}^{N}(1-b_i z)}{(1-az)},$$

then the density h of the exponential functional I for the compound Poisson process associated to X_1 is given by

$$h(x) = \sum_{n=0}^{+\infty} r_n \beta^{-n} e^{-\beta^{-n} x}, \qquad x \geq 0,$$

with, for $m, n \in \mathbb{N}$,

$$C_m = \sum_{m_1 + \cdots + m_N = m} \prod_{i=1}^{N} (-1)^{m_i} \beta^{m_i(m_i+1)/2} \frac{b_i^{m_i}}{(\beta; \beta)_{m_i}},$$

and

$$r_n = \frac{1}{\prod_{k=1}^{N}(b_k; \beta)_\infty}$$

$$\times \begin{cases} \displaystyle\sum_{m=0}^{n} C_m \frac{(a\beta^{m-n+1}; \beta)_\infty}{(\beta, \beta)_\infty (1/\beta; 1/\beta)_{n-m}}, & a \notin \{\beta^p : p \geq 1\}, \\[2em] \displaystyle\sum_{m=\max(n-p,0)}^{n} C_m \frac{1}{(1/\beta; 1/\beta)_{n-m}(\beta; \beta)_{m+p-n}}, & a = \beta^p. \end{cases}$$

PROOF.   The method is similar to the one used in the last example, (3.16) gives the relationship

$$\prod_{k=1}^{N}(b_k\beta; \beta)_\infty \mathbb{E}(e^{-\lambda I})$$

$$= \sum_{n=0}^{+\infty}(a\beta; \beta)_n\frac{(-\lambda)^n}{(\beta; \beta)_n}\prod_{i=1}^{N}(b_i\beta^{n+1}; \beta)_\infty$$

$$= \sum_{n=0}^{+\infty}(a\beta; \beta)_n\frac{(-\lambda)^n}{(\beta; \beta)_n}\sum_{(m_i)\in\mathbb{N}^N}\prod_{i=1}^{N}(-1)^{m_i}\beta^{m_i(m_i-1)/2}\frac{b_i^{m_i}\beta^{(n+1)m_i}}{(\beta; \beta)_{m_i}},$$

then

$$\prod_{k=1}^{N}(b_k\beta; \beta)_\infty \mathbb{E}(e^{-\lambda I})$$

$$= \sum_{n=0}^{+\infty}(a\beta; \beta)_n\frac{(-\lambda)^n}{(\beta; \beta)_n}\sum_{m=0}^{+\infty}C_m\beta^{nm}$$

$$= \sum_{m=0}^{+\infty}C_m\sum_{n=0}^{+\infty}(a\beta; \beta)_n\frac{(-\beta^m\lambda)^n}{(\beta; \beta)_n} = \sum_{m=0}^{+\infty}C_m\frac{(-a\beta^{m+1}\lambda; \beta)_\infty}{(-\beta^m\lambda; \beta)_\infty}$$

by the q-Binomial theorem. For $n \in \mathbb{N}$ the expression of the residue of the Laplace transform of I at $-\beta^{-n}$ is then easy to obtain. The proposition is proved.   □

4. The throughput of the AIMD algorithm.   In this section the study of the AIMD algorithm is completed. The variable X_1 is assumed to be integer valued and greater than 1. In the model considered in this paper the loss rate of packets is of the order $\alpha\mathbb{E}(X_1)$. Recall the definition of the throughput of an AIMD algorithm.

DEFINITION 10.   The throughput of the algorithm is defined as the limit

$$\rho^\alpha = \lim_{n\to+\infty}\frac{1}{n}\sum_{k=1}^{n}W_k^\alpha = \mathbb{E}(W_\infty^\alpha).$$

This definition assumes that the round trip time (RTT) is taken equal to 1. Recall that basically, RTT is the time interval between the transmission of two windows. Due to the fact that the occupancy of the buffers of the routers vary, the packets experience variable delays along their path. This implies, in particular, that the RTT will vary too, as it will be seen in a following discussion (Section 5), assuming that the RTT constant is by no means restrictive.

Using the embedded Markov chain (V_n^α), the throughput can also be written as

$$\rho^\alpha = \frac{\mathbb{E}\left(\sum_{k=0}^{G_{V_\infty^\alpha}^\alpha - 1}(V_\infty^\alpha + k)\right)}{\mathbb{E}(G_{V_\infty^\alpha}^\alpha)} = \frac{\mathbb{E}(2G_{V_\infty^\alpha}^\alpha V_\infty^\alpha + (G_{V_\infty^\alpha}^\alpha)^2)}{2\mathbb{E}(G_{V_\infty^\alpha}^\alpha)} - \frac{1}{2},$$

by multiplying this identity by the square root of the loss rate of packets, that is, by $\sqrt{\alpha \mathbb{E}(X_1)}$, Theorem 2 shows that the convergence

$$(4.1) \qquad \overline{\rho}_{X_1} \stackrel{\text{def.}}{=} \lim_{\alpha \to 0} \sqrt{\alpha \mathbb{E}(X_1)}\rho^\alpha = \sqrt{\mathbb{E}(X_1)}\frac{\mathbb{E}(2\overline{G}_{\overline{V}_\infty}\overline{V}_\infty + \overline{G}_{\overline{V}_\infty}^2)}{2\mathbb{E}(\overline{G}_{\overline{V}_\infty})}$$

holds. Since, by (2.6),

$$(\overline{V}_\infty + \overline{G}_{\overline{V}_\infty})^2 \stackrel{\text{dist.}}{=} \overline{V}_\infty^2 + 2E_1,$$

one gets the relationship

$$\mathbb{E}(2\overline{G}_{\overline{V}_\infty}\overline{V}_\infty + \overline{G}_{\overline{V}_\infty}^2) = \mathbb{E}((\overline{V}_\infty + \overline{G}_{\overline{V}_\infty})^2 - \overline{V}_\infty^2) = 2.$$

Equation (2.4) at equilibrium gives directly

$$\mathbb{E}(\overline{G}_{\overline{V}_\infty}) = \frac{1 - \mathbb{E}(\delta^{X_1})}{\mathbb{E}(\delta^{X_1})}\mathbb{E}(\overline{V}_\infty).$$

These last identities show that (4.1) can be rewritten as

$$(4.2) \qquad \overline{\rho}_{X_1} = \frac{\sqrt{\mathbb{E}(X_1)}\mathbb{E}(\delta^{X_1})}{(1 - \mathbb{E}(\delta^{X_1}))\mathbb{E}(\overline{V}_\infty)}.$$

The next proposition gives an explicit formula for the asymptotic throughput.

THEOREM 11. *The asymptotic throughput of an AIMD algorithm with multiplicative decrease factor δ in a correlated loss model associated to the random variable X_1 is given by*

$$(4.3) \qquad \overline{\rho}_{X_1} = \lim_{\alpha \to 0} \sqrt{\alpha \mathbb{E}(X_1)}\rho^\alpha = \sqrt{\frac{2\mathbb{E}(X_1)}{\pi}}\prod_{n=1}^{+\infty}\frac{1 - \mathbb{E}(\delta^{2nX_1})}{1 - \mathbb{E}(\delta^{(2n-1)X_1})}.$$

PROOF. According to the remark at the beginning of Section 3, the variable I, solution of (3.1) with $\beta = \delta^2$, and the variable $\overline{V}_\infty^2/2 + E_1$ (E_1 is exponentially distributed with parameter 1 and independent of $\overline{V}_\infty$) have the same distribution. Equation (2.5) gives the identity in distribution

$$\overline{V}_\infty^2/2 \stackrel{\text{dist.}}{=} \delta^{2X_1}(\overline{V}_\infty^2/2 + E_1) \stackrel{\text{dist.}}{=} \delta^{2X_1}I,$$

therefore $\mathbb{E}(\overline{V}_\infty) = \sqrt{2}\mathbb{E}(\delta^{X_1})E(\sqrt{I})$. Formula (3.12) yields

$$E(\sqrt{I}) = \Gamma(3/2)\prod_{n=1}^{+\infty}\frac{1 - \mathbb{E}(\delta^{(1+2n)X_1})}{1 - \mathbb{E}(\delta^{2nX_1})}.$$

Since $\Gamma(3/2) = \sqrt{\pi}/2$, (4.2) gives the desired formula. $\quad\square$

The impact of the correlation of the loss process.   The sensitivity of $\overline{\rho}_X$ with respect to the variance of X is now investigated. The goal is to compare models with the same loss rate [$\alpha\mathbb{E}(X_1)$ for the model considered up to now]. For this purpose the definitions of stochastic order and concave order are recalled. See [24] for the basic definitions and results on stochastic orderings.

DEFINITION 12.   The order relationships $\leq_{\text{st}}$ and $\leq_{\text{cv}}$ are defined as follows, for two random variables X and Y on $\mathbb{R}$:

1. The inequality $X \leq_{\text{st}} Y$ holds when

$$\mathbb{E}(f(X)) \leq \mathbb{E}(f(Y))$$

is true for any nondecreasing function on $\mathbb{R}$. Equivalently, $X \leq_{\text{st}} Y$ if and only if the inequality $\mathbb{P}(X \geq a) \leq \mathbb{P}(Y \geq a)$ holds for any $a \in \mathbb{R}$.
2. The inequality $X \leq_{\text{cv}} Y$ holds when

$$\mathbb{E}(f(X)) \leq \mathbb{E}(f(Y))$$

is true for any nondecreasing concave function on $\mathbb{R}$. Equivalently, $X \leq_{\text{cv}} Y$ if and only if the inequality $\mathbb{E}((a - Y)^+) \leq \mathbb{E}((a - X)^+)$ holds for any $a \in \mathbb{R}$.

If X and X' (resp. Y and Y') are independent real random variables such that $X \leq_{\text{cv}} Y$ and $X' \leq_{\text{cv}} Y'$, then $X + X' \leq_{\text{cv}} Y + Y'$. Indeed, for $a \in \mathbb{R}$,

$$\mathbb{E}\big((a - (Y + Y'))^+\big)$$

$$= \int_{\mathbb{R}} \mathbb{E}\big((a - y - Y')^+\big)\mathbb{P}(Y \in dy)$$

$$= \int_{\mathbb{R}} \mathbb{E}\big((a - y - X')^+\big)\mathbb{P}(Y \in dy) = \int_{\mathbb{R}} \mathbb{E}\big((a - x' - Y)^+\big)\mathbb{P}(X' \in dx')$$

$$\leq \int_{\mathbb{R}} \mathbb{E}\big((a - x' - X)^+\big)\mathbb{P}(X' \in dx') = \mathbb{E}\big((a - (X + X'))^+\big).$$

In the same way, under the same independence assumptions if X, X', Y and Y' are such that $X \leq_{\text{st}} Y$ and $X' \leq_{\text{st}} Y'$, then $X + X' \leq_{\text{st}} Y + Y'$.

If X and Y are random variables such that $X \leq_{\text{st}} Y$, it is possible to construct a common probability space for two random variables X' and Y' having respectively the same distribution as X and Y, and such that the inequality $X' \leq_{\text{st}} Y'$ holds almost surely. For the nonrenormalized loss rate $\overline{\rho}_Z/\sqrt{\mathbb{E}(Z)}$, the inequality

$$\frac{\overline{\rho}_X}{\sqrt{\mathbb{E}(X)}} \geq \frac{\overline{\rho}_Y}{\sqrt{\mathbb{E}(Y)}}$$

should hold since the model with X experiences less losses than the model with Y. From (4.3), this is not obvious at all. In this part the expression for the throughput is rewritten in a more convenient form to compare several loss processes.

PROPOSITION 13. *The asymptotic throughput $\overline{\rho}_{X_1}$ can be written as*

$$(4.4) \qquad \overline{\rho}_{X_1} = \sqrt{\frac{2\mathbb{E}(X_1)}{\pi}} \exp\left(\sum_{k=1}^{+\infty} \frac{1}{k}\mathbb{E}\left(\frac{1}{1+\delta^{-S_k}}\right)\right),$$

where $(S_n) = (X_1 + \cdots + X_n)$ is the random walk associated to the i.i.d. sequence (X_n).

PROOF. From (4.3), one gets (recall that $X_1 \geq 1$),

$$\log\left(\sqrt{\frac{\pi}{2\mathbb{E}(X_1)}}\,\overline{\rho}_{X_1}\right) = \sum_{n=1}^{+\infty} (-1)^n \log\left(1 - \mathbb{E}(\delta^{nX_1})\right)$$

$$= -\sum_{n=1}^{+\infty}\sum_{k=1}^{+\infty}(-1)^n \frac{1}{k}\left(\mathbb{E}(\delta^{nX_1})\right)^k = -\sum_{k=1}^{+\infty}\sum_{n=1}^{+\infty}\frac{1}{k}(-1)^n \mathbb{E}(\delta^{nS_k})$$

$$= \sum_{k=1}^{+\infty}\frac{1}{k}\mathbb{E}\left(\frac{\delta^{S_k}}{1+\delta^{S_k}}\right),$$

and (4.4) is established. $\square$

If $X \leq_{\mathrm{st}} Y$, the same property holds for the associated random walks (S_n^X) and (S_n^Y), that is, for any $n \geq 1$, $S_n^X \leq_{\mathrm{st}} S_n^Y$, thus (4.4) shows directly that $\overline{\rho}_X/\sqrt{\mathbb{E}(X)} \geq \overline{\rho}_Y/\sqrt{\mathbb{E}(Y)}$ as expected. The following proposition gives a stronger result in this domain.

PROPOSITION 14. *The asymptotic throughput $Z \to \overline{\rho}_Z/\sqrt{\mathbb{E}(Z)}$ is a non-increasing function for the concave order, that is, if X and Y are random variables*

$$(4.5) \qquad X \leq_{\mathrm{cv}} Y \quad \text{implies} \quad \frac{\overline{\rho}_X}{\sqrt{\mathbb{E}(X)}} \geq \frac{\overline{\rho}_Y}{\sqrt{\mathbb{E}(Y)}}.$$

In particular, when $\mathbb{E}(X) = \mathbb{E}(Y)$, $X \leq_{\mathrm{cv}} Y$ implies $\overline{\rho}_X \geq \overline{\rho}_Y$.

PROOF. If (S_n^X) and (S_n^Y) are random walks associated to the variables X and Y, by induction, with the help remark below Definition 12, it is easily seen that for $n \geq 1$, $S_n^X \leq_{\mathrm{cv}} S_n^Y$.

The function $a \to 1/(\delta^a + 1)$ being nondecreasing and concave on $\mathbb{R}_+$, one gets that for $n \geq 1$,

$$\mathbb{E}\left(\frac{1}{\delta^{S_n^X}+1}\right) \leq \mathbb{E}\left(\frac{1}{\delta^{S_n^Y}+1}\right), \qquad \mathbb{E}\left(\frac{1}{\delta^{-S_n^X}+1}\right) \geq \mathbb{E}\left(\frac{1}{\delta^{-S_n^Y}+1}\right).$$

By (4.4), the last inequality implies that $\overline{\rho}_X/\sqrt{\mathbb{E}(X)} \geq \overline{\rho}_Y/\sqrt{\mathbb{E}(Y)}$. This completes the proof. $\square$

The above proposition suggests the greater is the variance of X, the better is the asymptotic throughput. Jensen's inequality gives that for any concave function f on $\mathbb{R}_+$,

$$\mathbb{E}(f(X)) \leq f(\mathbb{E}(X)),$$

hence $\mathbb{E}(X) \geq_{\mathrm{cv}} X$. This implies, in particular, that $\overline{\rho}_X \geq \overline{\rho}_{\mathbb{E}(X)}$ where, for $t > 0$, $\overline{\rho}_t$ denotes the asymptotic throughput for the random variable constant equal to t. In other words, the asymptotic throughput with the loss process associated to X is greater than the throughput of an uncorrelated model but with a multiplicative decay $\delta^{\mathbb{E}(X)}$.

PROPOSITION 15.   *For any integer valued random variable X,*

$$(4.6) \qquad \overline{\rho}_X \geq \overline{\rho}_{\mathbb{E}(X)} = \sqrt{\frac{2\mathbb{E}(X)}{\pi}} \prod_{n=1}^{+\infty} \frac{1 - \delta^{2n\mathbb{E}(X)}}{1 - \delta^{(2n-1)\mathbb{E}(X)}}.$$

The function $t \to \overline{\rho}_t$ is nondecreasing, in particular, for any random variable $X \geq 1$,

$$(4.7) \qquad \overline{\rho}_X \geq \overline{\rho}_1.$$

According to (4.7), (4.3) with $X_1 \equiv 1$ proved in [8] is thus a lower bound for the real throughput when the packet losses are correlated. Equation (4.7) of the above proposition shows that choosing an uncorrelated loss process underestimates the real performance. A possible intuitive explanation of this phenomenon is the following: when there are x losses in some small time interval, the congestion window size is basically reduced by a factor δ^x. If x is not too small then, due to the exponential decay, for any $y \geq x$ the quantities δ^x or δ^y are both very small. Hence it is better to have a very large variability in the loss process: large number of losses locally but very rare.

PROOF OF PROPOSITION 15.   Only the nondecreasing property of $t \to \overline{\rho}_t$ has to be proved. Taking $g(t) = \log \overline{\rho}_t - \log(2/\pi)/2$ and using (4.4),

$$g(x) = \frac{1}{2} \log(x) + \sum_{n=1}^{+\infty} \frac{1}{n} \frac{1}{1 + \delta^{-nx}},$$

$$g'(x) = \frac{1}{2x} + \sum_{n=1}^{+\infty} \log(\delta) \frac{\delta^{nx}}{(1 + \delta^{nx})^2}.$$

A glance at the right-hand side of the last inequality reveals that to prove the property it is sufficient to show that the relationship

$$x \sum_{n=1}^{+\infty} \frac{e^{-nx}}{(1 + e^{-nx})^2} \leq \frac{1}{2}$$

holds for $x \geq 0$. Since

$$\sum_{n=1}^{+\infty} \frac{e^{-nx}}{(1+e^{-nx})^2} \leq \int_0^{+\infty} \frac{e^{-xy}}{(1+e^{-xy})^2} \, dy = \frac{1}{2x},$$

this is clearly true. The proposition is therefore proved. $\square$

As a consequence of Proposition 14 the asymptotic throughput is a nondecreasing functional, with respect to the concave order, for random variables X with the same mean value. When the mean values are different, the comparison turns out to be more difficult. Relationship (4.7) is an example of such a comparison for deterministic variables. This part is concluded with a simple example, when X is geometrically distributed, where this comparison is also possible. This kind of distribution also has another advantage: Since the number of local losses is believed to be sharply concentrated near small values (see [22]), the geometric distribution is a good candidate to describe the loss process.

PROPOSITION 16. *If, for $p \in [0, 1[$, G_p is a shifted geometrically distributed random variable with parameter p, that is, $\mathbb{P}(G_p = n) = p^{n-1}(1 - p)$ for $n \geq 1$, the function*

$$p \to \overline{\rho}_{G_p} = \sqrt{\frac{2}{\pi(1-p)}} \prod_{n=1}^{+\infty} \frac{1 - p\delta^{2n-1}}{1 - p\delta^{2n}} \frac{1 - \delta^{2n}}{1 - \delta^{2n-1}}$$

is convex and nondecreasing.

The case $p = 0$ corresponds to the uncorrelated case considered in [8]. Notice this is still the worst case for the asymptotic throughput. Recall that all these models have the same loss rate but with a variability increasing with p. In this case the mean value $\mathbb{E}(G_p)$ is no constant with p.

PROOF OF PROPOSITION 16. Equation (4.3) gives the relationship

$$\sqrt{\frac{\pi}{2}} \overline{\rho}_{G_p} = \frac{1}{\sqrt{1-p}} \prod_{n=1}^{+\infty} \frac{1 - p\delta^{2n-1}}{1 - p\delta^{2n}} \prod_{n=1}^{+\infty} \frac{1 - \delta^{2n}}{1 - \delta^{2n-1}},$$

thus,

$$\log\left(\frac{\overline{\rho}_{G_p}}{\overline{\rho}_{G_0}}\right) = -\frac{1}{2}\log(1 - p) - \sum_{n=1}^{+\infty}(-1)^n \log(1 - p\delta^n)$$

$$= \sum_{k=1}^{+\infty} \frac{1}{2k} p^k + \sum_{n=1}^{+\infty}\sum_{k=1}^{+\infty}(-1)^n \frac{p^k}{k}\delta^{nk} = \sum_{k=1}^{+\infty} \frac{p^k}{k}\left(\frac{1}{2} - \frac{\delta^k}{1 + \delta^k}\right)$$

$$= \sum_{k=1}^{+\infty} \frac{p^k}{k} \frac{1 - \delta^k}{2(1 + \delta^k)}.$$

The proposition is proved. □

5. On the accuracy of the stochastic model describing a TCP connection.
In this section several aspects of the TCP protocol not explicitly considered in the
stochastic model analyzed in this paper are discussed.

5.1. *Finite maximal congestion window size.* For the moment it has been
assumed that the sequence (W_n^α) can be increased without any bound. In practice,
the congestion window size is blocked as soon it has reached a maximal
value w_{max}^α. In [8] for independent packet losses, the stationary behavior of the
asymptotic sequence $(\overline{V}_n)$ is described when

$$w_{max}^\alpha \sim \overline{w}_{max}/\sqrt{\alpha}.$$

For the present loss model, a similar analysis can also be done. The corresponding
asymptotic sequence $(\overline{V}_n)$ satisfies the relation

$$\overline{V}_{n+1}^2 \overset{\text{dist.}}{=} \delta^{2X_n} \min(\overline{V}_n^2 + 2E_n, \overline{w}_{max}), \qquad n \geq 1,$$

where (X_n) and (E_n) are i.i.d. independent sequences and E_1 is exponentially
distributed with parameter 1. This sequence converges in distribution to a random
variable $\overline{V}_\infty$ such that

$$\overline{V}_\infty = \sqrt{\inf_{n \geq 0}\left(\delta^{2S_n}\overline{w}_{max} + 2\sum_{i=1}^{n}\delta^{2S_i}E_i\right)},$$

where (S_n) is the random walk associated to (X_n).

5.2. *Timeouts.* In the model considered here only losses that can be handled
by the Fast Recovery Algorithm have been considered (see [15]). When three
consecutive packets are lost or when the timeout for a packet has elapsed, the
congestion window size W is set to 1. In our limiting process this amounts to
set $\overline{W}$ to 0. Thus, the evolution equation

$$\overline{V}_1^2 \overset{\text{dist.}}{=} \delta^{2X_1}(\overline{V}_0^2 + 2E_1)$$

is still valid provided that the value $+\infty$ is allowed for X_1.

The quantity $\mathbb{P}(X_1 = +\infty)$ is then interpreted as the probability of a timeout or
of three consecutive losses. Formula (4.3) for the asymptotic throughput becomes

$$\overline{\rho}_{X_1} = \sqrt{\frac{2(q\mathbb{E}(X_1|X_1 < +\infty) + (1-q))}{\pi}}\prod_{n=1}^{+\infty}\frac{1 - q\mathbb{E}(\delta^{2nX_1}|X_1 < +\infty)}{1 - q\mathbb{E}(\delta^{(2n-1)X_1}|X_1 < +\infty)},$$

where $q = \mathbb{P}(X_1 < +\infty)$ the probability of a timeout is $1 - q$. In the simple
uncorrelated case, $\mathbb{P}(X_1 = 1) = q = 1 - \mathbb{P}(X_1 = +\infty)$, the above formula is

$$\overline{\rho}_1 = \sqrt{\frac{2}{\pi}}\prod_{n=1}^{+\infty}\frac{1 - q\delta^{2n}}{1 - q\delta^{2n-1}} = \sqrt{\frac{2}{\pi}}\exp\left(\sum_{k=1}^{+\infty}\frac{q^k}{k}\frac{\delta^k}{1 + \delta^k}\right),$$

where the last identity is proved by using the same arguments as in the proof of Proposition 16. For TCP ($\delta = 1/2$), as q varies from 1 to 0, $\overline{\rho}_1$ decreases from 1.309 to 0.798.

5.3. *Slow start phase.* In the present implementations of TCP, if an isolated loss (i.e., not within a group of three consecutive losses as before) occurs when the window size is $W = w$, the algorithm *Slow Start* is then used (see [23]); it works as follows. A quantity called *Slow Start Threshold* T_{ss} is fixed to $\lfloor w/2 \rfloor$ and the congestion window size is set to 1. The congestion window size is then doubled after each RTT as long as its value has not reached T_{ss}:

$$(5.1) \qquad W \to \begin{cases} 2W, & \text{if no loss occurs among the } W \text{ packets,} \\ 1, & \text{otherwise,} \end{cases}$$

when W is greater than T_{ss}, the AIMD algorithm is then used. Ferguson [10] analyses a related stochastic model of the slow start algorithm. In the probabilistic model investigated here (and also in [8]), this algorithm of the TCP protocol is taken into account.

In the setting of the paper this algorithm can be included without changing the results obtained so far. Recall, Proposition 4, that the chain (W_n^α) is properly renormalized as

$$(5.2) \qquad \left(\sqrt{\alpha} W_{\lfloor t/\sqrt{\alpha} \rfloor}^\alpha \right)$$

to get the asymptotic Markov process $(\overline{W}(t))$. To show that the slow start algorithm can be neglected, it is sufficient to show that if $W_0^\alpha = 1$ and the transitions (5.1) are used, the mean time T_x to reach the level $x/\sqrt{\alpha}$ is $o(1/\sqrt{\alpha})$. In other words, the time necessary to reach a slow start threshold of the order $x/\sqrt{\alpha}$ is negligible in the time scale defined by (5.2), therefore the slow start period vanishes because of the time scale.

PROPOSITION 17. *If $W_0^\alpha = \lfloor x/\sqrt{\alpha} \rfloor$, (W_n^α) is a TCP session starting after a loss and T^α is the first index n when the congestion avoidance algorithm is used, then*

$$\lim_{\alpha \to 0} \mathbb{P}\left(\frac{T^\alpha}{-\log_2 \sqrt{\alpha}} \leq 1 \right) = 1.$$

The variable T^α is of the order $-\log_2 \sqrt{\alpha}$, hence the interval $\{0, 1, \ldots, T^\alpha\}$ (where the slow start algorithm is used) vanishes under the scaling (5.2), $t \to \lfloor t/\sqrt{\alpha} \rfloor$, consequently so does the slow start algorithm in the stochastic model.

PROOF OF PROPOSITION 17. Since the distribution of the size of a group of losses is independent of α, it can be assumed that the initial loss is the last loss of a group. (Recall that time is shrunk by $1/\sqrt{\alpha}$.) The next loss will thus occur

as in the independent loss model, where each packet has a probability $\exp(-\alpha)$ of being lost. If no loss occurs during the first $\log_2\lfloor x/\sqrt{\alpha}\rfloor$ steps, then necessarily $T^\alpha \leq \log_2\lfloor x/\sqrt{\alpha}\rfloor$, therefore

$$\mathbb{P}(T^\alpha \leq \log_2\lfloor x/\sqrt{\alpha}\rfloor) \geq \prod_{i=1}^{\log_2\lfloor x/\sqrt{\alpha}\rfloor} \exp(-\alpha 2^i) = \exp(-\alpha(\lfloor x/\sqrt{\alpha}\rfloor - 1)).$$

Since this last expression is converging to 1 as α tends to 0, the proposition is proved. $\square$

REMARK. The slow start algorithm does not play any role in this paper because only the transfer of an infinite file is considered. The transient periods where the algorithm recovers from a loss are negligible from this point of view. The problem is entirely different when "small transfers" (less than ten packets say) are considered. For these connections, the reverse situation prevails and they are finished before the congestion avoidance algorithm starts.

5.4. *Variable RTTs.* In Section 4 devoted to the asymptotic throughput of the TCP connection, the round trip times have been assumed constant. In practice this is not the case since packets experienced delays in the buffers of the various routers along their paths.

If, for $n \in \mathbb{N}$, R_n is the delay experienced by the packets of the nth round trip, the random variables W_n^α and R_n are correlated random variables. The average throughput after the nth round trip is

$$\sum_{i=1}^{n} W_i^\alpha \bigg/ \sum_{i=1}^{n} R_i^\alpha.$$

If we assume that, asymptotically, the sequence (R_n) is stationary, the ergodic theorem shows that, almost surely,

$$\lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^{n} W_i^\alpha = \mathbb{E}(W_\infty^\alpha) \quad \text{and} \quad \lim_{n\to+\infty} \frac{1}{n}\sum_{i=1}^{n} R_i = \mathbb{E}(R_\infty);$$

the asymptotic throughput is therefore $\mathbb{E}(W_\infty)/\mathbb{E}(R_\infty)$. Notice that the dependence between the two sequences (W_n^α) and (R_n) does not play any role for this result. Hence, up to the constant $1/\mathbb{E}(R_\infty)$, and under a mild assumption on the stationarity of (R_n), even when the RTT are variable, Theorem 11 gives the right expression for the asymptotic throughput for a long TCP connection.

REMARK. In [5], it is shown that a self-similar process plays an important role for the exponential functionals. Bertoin, Biane and Yor [4] give a nice probabilistic representation of the distribution of this process when the Lévy process is Poisson. In spite of the occurrence this self-similar process is quite appealing, it does not seem that it has an interpretation in the stochastic model of TCP.

APPENDIX

Elementary results concerning q-calculus are recalled in this section. See [3] and [17] for a quick introduction and Gasper and Rahman [13] for a more advanced presentation. Recall that for $x \in \mathbb{C}$, $q \in [0, 1[$ and $n \in \mathbb{N} \cup \{+\infty\}$,

$$(x; q)_n = (1 - x)(1 - xq)(1 - xq^2) \cdots (1 - xq^{n-1})$$

for $n \geq 1$ and $(a; q)_0 = 1$. For $n < +\infty$, the quantity $(q; q)_n$ is a generalized factorial in the sense that

$$\lim_{q \nearrow 1} \frac{(q; q)_n}{(1 - q)^n} = n!.$$

Roughly speaking, q-calculus is ordinary calculus but with factorials replaced by these generalized factorials.

The q-hypergeometric functions $_r\phi_s$ are defined by

(A.1)
$$_r\phi_s \begin{pmatrix} a_1, \ldots, a_r \\ b_1, \ldots, b_s \end{pmatrix}; q, z \end{pmatrix}$$
$$= \sum_{m=0}^{\infty} \frac{(a_1; q)_m \cdots (a_r; q)_m}{(b_1; q)_m \cdots (b_s; q)_m} ((-1)^m q^{m(m-1)/2})^{1+s-r} \frac{z^m}{(q; q)_m}$$

for r and $s \in \mathbb{N}$. The q-hypergeometric functions $_r\phi_s$ are generalized versions of the classical higher-order hypergeometric function $_rF_s$,

$$\lim_{q \nearrow 1} {}_r\phi_s \begin{pmatrix} a_1, \ldots, a_r \\ b_1, \ldots, b_s \end{pmatrix}; q, (q-1)^{1+s-r} z \end{pmatrix} = {}_rF_s \begin{pmatrix} a_1, \ldots, a_r \\ b_1, \ldots, b_s \end{pmatrix}; z \end{pmatrix}.$$

THEOREM 18 (q-Binomial theorem). *For $|x| < 1$ and $|q| < 1$,*

(A.2)
$$_1\phi_0(a; q, x) = \sum_{k=0}^{+\infty} (a; q)_k \frac{x^k}{(q; q)_k} = \frac{(ax; q)_\infty}{(x; q)_\infty}.$$

Euler's formulas are direct consequences of (A.2),

(A.3)
$$\sum_{k=0}^{+\infty} (-1)^k q^{k(k-1)/2} \frac{x^k}{(q; q)_k} = (x; q)_\infty,$$

(A.4)
$$\sum_{k=0}^{+\infty} \frac{x^k}{(q; q)_k} = \frac{1}{(x; q)_\infty}.$$

The first one is obtained by taking $a = 0$ and for the second one, a (resp. x) is replaced by $1/a$ (resp. ax), and then a is set to 0.

REFERENCES

[1] ALDOUS, D. (1989). *Probability Approximations via the Poisson Clumping Heuristic*. Springer, New York.

[2] ALTMAN, E., AVRACHENKO, K., BARAKAT, C. and NUÑEZ-QUEIJA, R. (2001). State dependent $M/G/1$ type queueing analysis for congestion control in data networks. In *Infocom 2001*. IEEE, New York.

[3] ANDREWS, G. E., ASKEY, R. and ROY, R. (1999). *Special Functions*. Cambridge Univ. Press.

[4] BERTOIN, J., BIANE, P. and YOR, M. (2002). Poissonian exponential functionals, q-series, q-integrals and the moment problem for log-normal distributions. Technical Report PMA-705, Laboratoire de Probabilités, Univ. Paris 6.

[5] BERTOIN, J. and YOR, M. (2001). On the entire moments of self-similar Markov processes and exponential functionals of Lévy processes. Technical Report PMA-702, Laboratoire de Probabilités, Univ. Paris 6.

[6] BOLOT, J. (1993). End-to-end packet delay and loss behavior in the Internet. In *ACM SIGCOMM'93* 289–298. ACM Press, New York.

[7] CARMONA, P., PETIT, F. and YOR, M. (1997). On the distribution and asymptotic results for exponential functionals of Lévy processes. In *Exponential Functionals and Principal Values Related to Brownian Motion* (M. Yor, ed.) 73–130. Rev. Mat. Iberoam., Madrid.

[8] DUMAS, V., GUILLEMIN, F. and ROBERT, P. (2002). A Markovian analysis of Additive-Increase Multiplicative-Decrease (AIMD) algorithms. *Adv. in Appl. Probab.* **34** 85–111.

[9] FELLER, W. (1971). *An Introduction to Probability Theory and Its Applications*, 2nd ed. Wiley, New York.

[10] FERGUSON, T. S. (1972). Lose a dollar or double your fortune. *Proc. Sixth Berkeley Symp. Math. Statist. Probab.* **3** 657–666. Univ. California Press, Berkeley.

[11] FLAJOLET, P., GOURDON, X. and DUMAS, P. (1995). Mellin transforms and asymptotics: Harmonic sums. *Theoret. Comput. Sci.* **144** 3–58.

[12] FLOYD, S. (1991). Connections with multiple congested gateways in packet-switched networks 1: One way traffic. *Computer Comm. Rev.* **21** 30–47.

[13] GASPER, G. and RAHMAN, M. (1990). *Basic Hypergeometric Series*. Cambridge Univ. Press.

[14] HARRISON, J. M. (1977). Ruin problems with compounding assets. *Stochastic Process. Appl.* **5** 67–79.

[15] JACOBSON, V. (1988). Congestion avoidance and control. In *SIGCOMM'88 Symposium*: *Communications Architectures and Protocols*. ACM, New York.

[16] KINGMAN, J. F. C. (1993). *Poisson Processes*. Clarendon Press, New York.

[17] KOORNWINDER, T. H. (1994). q-special functions, a tutorial. In *Representation of Lie Groups and Quantum Groups* (V. Baldoni and M. A. Picardello, eds.) 46–128. Longman, New York.

[18] MISRA, V., GONG, W. and TOWSLEY, D. (1999). Stochastic differential equation modeling and analysis of TCP windowsize behavior. In *Performance'99*.

[19] NEVEU, J. (1977). Processus ponctuels. *École d'Été de Probabilités de Saint-Flour VI. Lecture Notes in Math.* **598** 249–445. Springer, Berlin.

[20] OTT, T. J., KEMPERMAN, J. H. B. and MATHIS, M. (1996). The stationary behavior of ideal TCP congestion avoidance. Unpublished manuscript.

[21] PADHYE, J., FIROIU, V., TOWSLEY, D. and KUROSE, J. (1998). Modeling TCP throughput: A simple model and its empirical validation. In *ACM Sigcomm'98*. ACM, New York.

[22] PAXSON, V. (1999). End-to-end internet packet dynamics. *IEEE/ACM Trans. Networking* **7** 277–292.

[23] STEVENS, W. R. (1994). *TCP-IP Illustrated* **1**: *The Protocols*. Addison–Wesley, Reading, MA.

[24] STOYAN, D. (1983). *Comparison Methods for Queues and Other Stochastic Models*. Wiley, New York.

[25] YAJNIK, M., KUROSE, J. and TOWSLEY, D. (1995). Packet loss correlation in the MBone multicast network experimental measurements and Markov chain models. Technical Report UM-CS-1995-115, Univ. Massachusetts, Amherst.

[26] YOR, M. (1992). *Some Aspects of Brownian Motion. Part I. Some Special Functionals*. Birkhäuser, Basel.

[27] YOR, M. (1997). *Some Aspects of Brownian motion. Part II. Some Recent Martingale Problems*. Birkhäuser, Basel.

[28] YOR, M. (2001). *Exponential Functionals of Brownian Motion and Related Processes*. Springer, Berlin.

[29] ZHANG, Y., PAXSON, V. and SHENKER, S. (2000). The stationarity of internet path properties: Routing, loss, and throughput. Technical report ACIRI.

F. GUILLEMIN
FRANCE TELECOM R&D, DAC/CPN
2, AVENUE PIERRE MARZIN
22 300 LANNION
FRANCE
E-MAIL: Fabrice.Guillemin@francetelecom.com

P. ROBERT
INRIA, DOMAINE DE VOLUCEAU
B.P. 105
78153 LE CHESNAY CEDEX
FRANCE
E-MAIL: Philippe.Robert@inria.fr

B. ZWART
DEPARTMENT OF MATHEMATICS
  AND COMPUTER SCIENCE
EINDHOVEN UNIVERSITY OF TECHNOLOGY
HG 9.35
P.O. BOX 513
5600 MB EINDHOVEN
THE NETHERLANDS
E-MAIL: zwart@win.tue.nl