Retransmissions over Correlated Channels

Predrag R. Jelenković Evangelia D. Skiani Department of Electrical Engineering Columbia University, New York, NY 10027, USA Email: {predrag, valia}@ee.columbia.edu

ABSTRACT

Frequent failures characterize many existing communication networks, e.g. wireless ad-hoc networks, where retransmission-based failure recovery represents a primary approach for successful data delivery. Recent work has shown that retransmissions can cause power law delays and instabilities even if all traffic and network characteristics are super-exponential. While the prior studies have considered an independent channel model, in this paper we extend the analysis to the practically important dependent case. We use modulated processes, e.g. Markov modulated, to capture the channel dependencies. We study the number of retransmissions and delays when the hazard functions of the distributions of data sizes and channel statistics are proportional, conditionally on the channel state. Our results show that the tails of the retransmission and delay distributions are asymptotically insensitive to the channel correlations and are determined by the state that generates the lightest asymptotics. This insight is beneficial both for capacity planning and channel modeling since we do not need to account for the correlation details. However, these results may be overly optimistic when the best state is infrequent, since the effects of 'bad' states may be prevalent for sufficiently long to downgrade the expected performance.

1. INTRODUCTION

Recovery mechanisms are employed in almost all engineering systems that are prone to failures. Restarting the interrupted jobs from the beginning after a failure occurs is a straightforward and often used failure recovery mechanism. It was first shown in [3, 16] that processing light-tailed (e.g. exponential) jobs with 'restart' mechanism may result in power law delays. Furthermore, it was recognized in [9] that widely used retransmission failure recovery in communication networks is a form of 'restart' mechanism and, thus, can cause heavy-tailed (long) delays even if all data, channel and protocol characteristics are light-tailed. In addition, it was observed in [9] that retransmission phenomena can lead to zero throughput and system instabilities, and therefore need to be carefully considered for the design of fault tolerant communication systems, especially for wireless ad-hoc networks, where failures are typical.

More specifically, it has been shown that power law delays arise in different layers of the networking architecture, where retransmission-based protocols are used, e.g. ALOHA type protocols in MAC layer [10, 7], end-to-end acknowledgements in transport layer [8, 13] as well as in other layers [9]. Retransmission induced heavy tails represent a novel explanation for the broadly observed power law delays in communication networks. For other (non-retransmission) mechanisms that can give rise to heavy tails see [12] and the references therein.

Previous studies consider an i.i.d model that was first introduced in [3] and further studied in [1, 9, 13] to describe the channel dynamics. In practice, communication channels are highly correlated in the sense that they switch between states with different characteristics. We extend the previously studied independent model [9] to the dependent case where the availability periods depend on the channel state. In order to capture the channel dependencies, we introduce a modulating process, e.g. see [12], where the distributions of the channel availability periods depend on the current state of the channel.

The proposed model is as follows. Let $\{J_n\}_{n\geq 1}$ be a stationary and ergodic modulating process with finitely many states $\{1, 2, \ldots, K\}$. Now, let $\{A_n(k), n \geq 1, k = 1, \ldots, K\}$ be a family of independent random variables, independent of $\{J_n\}$, such that for fixed k, $\{A_n(k)\}_{n\geq 1}$ are identically distributed with $\bar{G}_k(x) = \mathbb{P}(A_1(k) > x)$. Then, we can construct a modulated process A_n such that $A_n := A_n(J_n)$ and $\mathbb{P}(A_n > x|J_n = k) = \bar{G}_k(x)$. At each available period A_n , we transmit a generic data unit of size L; if $A_n > L$, the transmission is successful, else we wait until the next period A_{n+1} to retransmit. We study the asymptotic properties of the distribution of the number of retransmissions N when

$$\log \mathbb{P}(L > x) \approx \alpha_k \log \mathbb{P}(A_n > x | J_n = k), \tag{1}$$

 $k = 1, \ldots, K$; see Subsection 1.1 for a more detailed description of the preceding model.

We show that when the channel is correlated, or less formally, when it alternates between different states, the tail asymptotics is determined by the properties of the 'best' channel state, e.g. the state that generates the lightest asymptotics in the corresponding independent channel model. Intuitively, as the channel switches between states, a large data unit is more likely to be transmitted when the channel is 'good'. Specifically, the 'best' availability periods correspond to the state with the largest α_k [as defined in (1)] among 1,..., K. Undoubtedly, this is an optimistic observation which further implies that instabilities (infinite ex-

 $^{^* {\}rm This}$ work is supported by the National Science Foundation grant number 0915784.

IFIP WG 7.3 Performance 2013, September 24-26, Vienna, Austria. Copyright is held by author/owner(s).

pected delay) can be eliminated as long as there exists at least one state with $\alpha > 1$.

From an engineering perspective, this optimistic best case scenario prediction and the apparent insensitivity to the structure of the channel correlations can be very promising in system analysis and design. The result implies that the initial i.i.d. model might be sufficient for modeling, and can also be extended to even more complex failure-prone networks. However, this is partially true as there are certain circumstances under which this claim underestimates the intricacies of the system.

Specifically, the lighter tail does not guarantee consistently good behavior for the entire body of the delay distribution. As discussed in [5, 6] in a different context, the delay distribution of bounded documents will always have a light exponential tail. However, the main body of the distribution can be a power law which may determine the performance in the relevant range of probabilities. Similarly here, the tail is determined by the lightest power law (largest value of α_k in (1)). However, when the corresponding channel state is rare, the main body of the retransmission distribution can be dominated by the heavier power laws resulting from the 'bad' states. Hence, the system performance may be much worse than predicted by the tail. Therefore, when the best case scenario is atypical, we need to pay closer attention to the channel correlations. We provide further discussions and some preliminary analysis of this situation in Section 4.

In engineering, our results can be used both for modeling and system design. In view of the observed insensitivity and state space reduction, it is likely that this analysis can be extended to more complex networks or multi-channel systems that are characterized by frequent failures and correlated states. The results may be applied in designing new protocols, or developing new fragmentation schemes [11, 14] specifically for correlated channels. A dynamic fragmentation technique (see [11]) is more likely to exhibit better performance in a channel with high variability. In addition, the explicit approximation presented in [6] could be combined with the analytical results of this paper in order to accurately estimate the optimal sizes of the packet fragments.

The paper is organized as follows. In the following Subsection 1.1, we formally describe the model and introduce the necessary definitions and notation, while in Section 2, we present our main theorems, on both the logarithmic and the exact scale. Next, Section 3 includes our simulation experiments that verify our analytic approximations. Last, in Section 4, we discuss the engineering implications of our results and provide some insight on the situation when the 'best case' scenario occurs rarely, while Section 5 contains some of the proofs.

1.1 Description of the Channel

In this section, we formally describe our model and provide necessary definitions and notation. Consider transmitting a generic data unit of random size L over a channel with failures. Without loss of generality, we assume that the channel is of unit capacity. The channel dynamics is modeled as follows. Let $\{J_n\}_{n\geq 1}$ be a stationary and ergodic modulating process with finitely many states $\{1, 2, \ldots, K\}$. Let $\pi_k = \mathbb{P}(J_n = k)$ denote the stationary probability that the process is in state $k, k = 1 \dots K$; let J be a generic random variable whose distribution is given by $\pi_k = \mathbb{P}(J = k)$.

Now, define a family of independent random variables

 $\{A_n(k), n \geq 1, k = 1, \ldots, K\}$, independent of the modulating process $\{J_n\}$. In addition, for fixed k, $\{A_n(k)\}_{n\geq 1}$ are identically distributed with $G_k(x) = \mathbb{P}(A_1(k) \leq x)$. Then, we construct a modulated process A_n such that $A_n := A_n(J_n)$ and $\mathbb{P}(A_n \leq x | J_n = k) = G_k(x)$. Note that the constructed process $\{A_n\}$ is also stationary and ergodic.

At each period of time that the channel becomes available, say A_i , we attempt to transmit a generic data unit of size L. If $A_i > L$, we say that the transmission is successful; otherwise, we wait until the next period A_{i+1} when the channel is available and attempt to retransmit the data from the beginning. A sketch of the model depicting the system is drawn in Figure 1.



Figure 1: Packets sent over a channel with failures

We are interested in computing the number of attempts N (retransmissions) that are required until L is successfully transmitted, which is formally defined as follows.

DEFINITION 1.1. The total number of retransmissions for a generic data unit of length L is defined as

$$N := \inf\{n : A_n > L\}$$

Moreover, we define the total transmission time.

DEFINITION 1.2. The total transmission time is defined as the total time until the data unit L is successfully delivered and is denoted as

$$T := \sum_{i=1}^{N-1} A_i + L.$$

We denote the conditional complementary cumulative distribution functions for $\{A_n\}_{n>1}$ and L, respectively, as

$$\bar{G}_k(x) := 1 - G_k(x) = \mathbb{P}(A_n > x | J_n = k)$$

and

$$\bar{F}(x) := \mathbb{P}[L > x].$$

Throughout the paper we assume that L and A are continuous (equivalently, $\overline{F}(x)$ and $\overline{G}_k(x)$ are absolutely continuous) and have infinite support, i.e. $\overline{G}_k(x) > 0$ and $\overline{F}(x) > 0$ for all $x \ge 0$. We use the following standard notations. For any two real functions a(t) and b(t) and fixed $t_0 \in \mathbb{R} \bigcup \{\infty\}$, we use $a(t) \sim b(t)$ as $t \to t_0$ to denote $\lim_{t\to t_0} a(t)/b(t) = 1$. Similarly, we say that $a(t) \gtrsim b(t)$ as $t \to t_0$ if $\lim_{t\to t_0} [a(t)/b(t)] \ge 1$; $a(t) \lesssim b(t)$ has a complementary definition.

2. MAIN RESULTS

In this section, we present our main analytical results. In Theorem 2.1, we characterize the tail distribution of the number of retransmissions on the logarithmic scale. In particular, under the assumption that the hazard functions of the data sizes and channel statistics are proportional, we show that the distribution exhibits a power law tail and the index is determined by the channel state with the longest availability periods, e.g. the maximum α_k as defined in (1). In other words, the asymptotics is the same as in the case of the uncorrelated channel with $\bar{G}(x) = \bar{G}_m(x)$, see Theorem 2 in [9], where m denotes the index of the 'best' state with maximum α_k .

Next, in Theorem 2.2, under more restrictive assumptions, we prove the result on the exact scale. We derive the exact asymptotic tail of the number of retransmissions N, which also depends on the steady state probability π_m of the 'best' state m. Last, Theorem 2.3 presents the logarithmic asymptotics of the total delay distribution.

Throughout the paper, we will use m to denote the index of the state with the largest α among all states $1, 2, \ldots, K$. This corresponds to the state m which dominates the tail distribution of N and is responsible for the lighter asymptote for large n. Without loss of generality, we assume that there is a unique m that achieves the maximum $\alpha_m = \max_k \alpha_k$, i.e. $\alpha_m > \max_{k \neq m} \alpha_k$. Otherwise, if there is more than one index that attains the maximum, we can merge the corresponding underlying states of the process $\{J_n\}_{n \geq 1}$ into a single one that uniquely achieves the maximum.

Furthermore, to simplify the notation, we sometimes omit the indices in A_n and J_n and simply write $\mathbb{P}(A_n > x | J_n = k) = \mathbb{P}(A > x | J = k)$.

THEOREM 2.1. Let $\{J_n\}_{n\geq 1}$ be a stationary and ergodic process that takes values on the positive integers k = 1...K. Assume that

$$\log \mathbb{P}(L > x) \sim \alpha_k \log \mathbb{P}(A > x | J = k) \text{ as } x \to \infty,$$

and $\mathbb{P}(|\sum_{i=1}^{n} 1\{J_i = k\} - \pi_k n| \ge \epsilon n) = O(n^{-(\alpha_m + \epsilon)/K})$, for positive ϵ and $\alpha_m := \max_{k=1...K} \alpha_k > 0$, then

$$\lim_{n \to \infty} \frac{\log \mathbb{P}(N > n)}{\log n} = -\alpha_m$$

REMARK 1. Note that the condition $\mathbb{P}(|\sum_{i=1}^{n} 1\{J_i = k\} - \pi_k n| \ge \epsilon n) = O(n^{-(\alpha_m + \epsilon)/K})$ is satisfied for a large class of modulating processes J_n , i.e. semi-Markov processes where the sojourn time distributions decay fast enough. For example, for the class of renewal Markov processes considered in Section 5 of [4], it can be shown that the condition holds if the sojourn time distribution is lighter than $O(n^{-(\alpha_m + \epsilon)/K+1})$.

Proof: By assumption, there exists $0 < \epsilon < 1$ such that for all $x \ge x_0$,

$$\bar{F}(x)^{\frac{1}{\alpha_k(1-\epsilon)}} \le \bar{G}_k(x) \le \bar{F}(x)^{\frac{1}{\alpha_k(1+\epsilon)}}, \qquad k = 1 \dots K.$$
(2)

Recall that $\{A_n\}_{n\geq 1}$ are conditionally independent given $\{J_n\}_{n\geq 1}$ and $\mathbb{P}(A_n > x | J_n = k) = \bar{G}_k(x)$. Note that $A_i(J_i)$ is independent of the past and future states of the modulating process $\{J_j\}_{j\neq i}$ given J_i . Let $N_n^k := \sum_{i=1}^n 1\{J_i = k\}$ be the number of times that $\{J_i = k\}$ in interval [1, n] and observe that $\sum_{k=1}^K N_n^k = n$.

First, we establish the *lower bound*. It is easy to see that

$$\mathbb{P}(N > n|L) = \mathbb{P}[L > A_1, L > A_2 \dots, L > A_n|L]$$

$$= \mathbb{E}_L[\mathbb{P}(L > A_j, 1 \le j \le n|J_1, \dots, J_n, L)]$$

$$= \mathbb{E}_L\left[\prod_{j=1}^n \mathbb{P}(L > A_j|J_j, L)\right]$$

$$= \mathbb{E}_L\left[\prod_{j=1}^n \prod_{k=1}^K \mathbb{P}(L > A_j|J_j = k, L)^{\mathbf{1}\{J_j = k\}}\right]$$

$$= \mathbb{E}_L\left[\prod_{j=1}^n \prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{\mathbf{1}\{J_j = k\}}\right]$$

$$= \mathbb{E}_L\left[\prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{N_n^k}\right], \quad (3)$$

where $\mathbb{E}_L[\cdot] = \mathbb{E}[\cdot|L]$. For the ergodic and stationary process $\{J_n\}_{n\geq 1}$, by the weak law of large numbers, it follows that

$$\mathbb{P}[N_n^k \le (1+\epsilon)\pi_k n] \to 1 \quad \text{as } n \to \infty,$$

for all k = 1...K. Thus, for any $\epsilon > 0$, we can choose n_0 , such that $\mathbb{P}[N_n^k \leq (1+\epsilon)\pi_k n] \geq (1-\epsilon)$, for all $n \geq n_0$ and k = 1...K. Therefore,

$$\mathbb{P}(N > n|L)$$

$$\geq \mathbb{E}_{L} \left[\prod_{k=1}^{K} \mathbb{P}(L > A|J = k, L)^{(1+\epsilon)\pi_{k}n} \mathbf{1} \left\{ N_{n}^{k} \leq (1+\epsilon)\pi_{k}n \right\} \right]$$

$$= \prod_{k=1}^{K} \left(\mathbb{P}(N_{n}^{k} \leq (1+\epsilon)\pi_{k}n)\mathbb{P}(L > A|J = k, L)^{(1+\epsilon)\pi_{k}n} \right)$$

$$\geq (1-\epsilon)^{K} \prod_{k=1}^{K} (1-\bar{G}_{k}(L))^{(1+\epsilon)\pi_{k}n}.$$

Now, using our main assumption (2) and the elementary inequality $1 - x \ge e^{-(1+\epsilon)x}$ for small x, we obtain

$$\mathbb{P}(N > n) = \mathbb{E}[\mathbb{P}(N > n|L)]$$

$$\geq (1 - \epsilon)^{K} \mathbb{E}\left[\prod_{k=1}^{K} \left(1 - \bar{F}(L)^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right)^{\pi_{k}n(1+\epsilon)} \mathbf{1}\{L \ge x_{0}\}\right]$$

$$\geq (1 - \epsilon)^{K} \mathbb{E}\left[\prod_{k=1}^{K} \exp\left(-\pi_{k}n(1+\epsilon)^{2}\bar{F}(L)^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right) \mathbf{1}\{L \ge x_{0}\}\right]$$

for x_0 as in (2). Next, observe that the expectation in the preceding expression is

$$\mathbb{E}\left[\exp\left(-\sum_{k=1}^{K}\pi_{k}n(1+\epsilon)^{2}U^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right)\mathbf{1}\{L\geq x_{0}\}\right]$$

$$\geq \mathbb{E}\left[\exp\left(-\sum_{k=1}^{K}\pi_{k}n(1+\epsilon)^{2}U^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right)\right]$$

$$-\exp\left(-\sum_{k=1}^{K}\pi_{k}n(1+\epsilon)^{2}\bar{F}(x_{0})^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right)$$

$$:=I_{1}-I_{0},$$
(4)

which follows from $\overline{F}(L) = U$, with U being uniformly distributed in (0, 1) by Proposition 2.1 in Chapter 10 of [15].

The first term in (4) is computed as

$$I_{1} = \int_{0}^{1} \exp\left(-\sum_{k=1}^{K} \pi_{k} n(1+\epsilon)^{2} u^{\frac{1}{\alpha_{k}(1+\epsilon)}}\right) du$$

$$\geq \int_{0}^{\epsilon} \exp\left(-n(1+\epsilon)^{2} u^{\frac{1}{\alpha_{m}(1+\epsilon)}} \pi_{m}$$

$$\cdot \left(1+\sum_{k=1, k\neq m}^{K} \frac{\pi_{k}}{\pi_{m}} u^{\frac{1}{\alpha_{k}(1+\epsilon)}-\frac{1}{\alpha_{m}(1+\epsilon)}}\right)\right) du$$

$$\geq \int_{0}^{\epsilon} \exp\left(-n(1+\epsilon)^{2}(1+\delta) u^{\frac{1}{\alpha_{m}(1+\epsilon)}} \pi_{m}\right) du,$$

where we observe that for ϵ small enough, $\sum_{k=1,k\neq m}^{K} (\pi_k/\pi_m) u^{1/(\alpha_k(1+\epsilon))-1/(\alpha_m(1+\epsilon))} \leq \delta$. Next, by change of variables $z = n\pi_m(1+\epsilon)^2(1+\delta)u^{1/(\alpha_m(1+\epsilon))}$, and setting $(1-\delta_{\epsilon})^{-1} = (1+\epsilon)^2(1+\delta)$, we have

$$I_{1} \geq \frac{(1-\delta_{\epsilon})^{\alpha_{m}(1+\epsilon)}\alpha_{m}(1+\epsilon)}{\pi_{m}^{\alpha_{m}(1+\epsilon)}n^{\alpha_{m}(1+\epsilon)}}$$
$$\cdot \int_{0}^{n\epsilon^{1/\alpha_{m}}(1-\delta_{\epsilon})\pi_{m}} e^{-z} z^{\alpha_{m}(1+\epsilon)-1} dz$$
$$\sim \frac{(1-\delta_{\epsilon})^{\alpha_{m}(1+\epsilon)}}{\pi_{m}^{\alpha_{m}(1+\epsilon)}n^{\alpha_{m}(1+\epsilon)}} \Gamma(\alpha_{m}(1+\epsilon)+1),$$

where we use the definition and the property of the Gamma function for large n.

Now, for $h_{\epsilon,\delta} = (1-\epsilon)^K (1-\delta_{\epsilon})^{\alpha_m(1+\epsilon)} \Gamma(\alpha_m(1+\epsilon) + 1)/\pi_m^{\alpha_m(1+\epsilon)}$, and, since x_0 is fixed, the second term in (4) is negligible, i.e. $I_0 \to 0$ exponentially fast as $n \to \infty$. Hence, taking the logarithm yields

$$\log \mathbb{P}(N > n) \ge \log h_{\epsilon,\delta} - \alpha_m (1 + 2\epsilon) \log n,$$

and by picking n_0 such that $\log h_{\epsilon,\delta} \geq -\alpha_m \epsilon \log n$, we have

$$\log \mathbb{P}(N > n) \ge -\alpha_m (1 + 3\epsilon) \log n.$$

After replacing ϵ with $\epsilon/3$, we derive

$$\frac{\log \mathbb{P}(N > n)}{\log n} \ge -(1 + \epsilon)\alpha_m.$$
(5)

The remainder of the proof of the *upper bound* follows similar arguments and is deferred to Section 5. \Box

Next, as briefly stated in the beginning of this section, we present our analytical approximation on the exact scale. In the following theorem, we need more restrictive assumptions and, in particular, we assume that the matching between the distribution of the data sizes and the channel characteristics is described by a regularly varying function of index α . A regularly varying function is defined as a function of the form $\ell(x)x^{\alpha}$, where $\ell(\cdot)$ is slowly varying.

DEFINITION 2.1. A function $\ell(x)$ is slowly varying if $\ell(x)/\ell(\lambda x) \to 1$ as $x \to \infty$ for any fixed $\lambda > 0$.

Typical examples of slowly varying functions are positive constants, logarithmic, e.g. $\log^{\beta} x, \beta \in R, \log(\log x)$, and $e^{(\log x)^{\gamma}}, 0 < \gamma < 1$. We assume that functions $\ell(x)$ are positive and bounded on finite intervals.

THEOREM 2.2. Let
$$m := \arg \max_{k=1...K} \alpha_k$$
. If $\bar{F}^{-1}(x) = \Phi_k(\bar{G}_k^{-1}(x))$, where $\Phi_k(x) = \ell(x)x^{\alpha_k}$, for all $x \ge 0, \alpha_k > 0$

 $\begin{array}{ll} 0,k = 1 \dots K, \ and \ \mathbb{P}(|\sum_{i=1}^{n} 1\{J_i = k\} - \pi_k n| \geq \epsilon n) = \\ O(n^{-(\alpha_m + \epsilon)/K}), \ \epsilon > 0, \ then \ as \ n \to \infty, \end{array}$

$$\mathbb{P}(N > n) \sim \frac{\Gamma(\alpha_m + 1)}{n^{\alpha_m} \ell(n) \pi_m^{\alpha_m}}.$$

Proof: The proof is deferred to Section 5.

Last, we prove the logarithmic asymptotics of the total transmission time, which is also a power law.

THEOREM 2.3. Under the same conditions as in Theorem 2.1 and $\mathbb{E}[A^{1+\theta}] < \infty$ for some $\theta > 0$, we obtain

$$\lim_{t \to \infty} \frac{\log \mathbb{P}(T > t)}{\log t} = -\alpha_m.$$

Proof: See Section 5.

3. SIMULATIONS

In this section, we illustrate the validity of our theoretical results with simulation experiments. In all of the experiments, we observed that the theoretical tail asymptotics is literally indistinguishable from the simulation. In the following examples, we present the simulation experiments resulting from 10^8 (or more) samples of $N_{i,1} \le i \le 10^8$. This number of samples was needed to ensure at least 100 occurrences in the lightest end of the tail that is presented in the figures ($N_i \ge n_{\max}$). This provides a good confidence interval given that we used simple two and three-state Markov chains to model $\{J_n\}$.



Figure 2: Example 1(a). Asymptotics of $\mathbb{P}(N > n)$ for a two-state channel.

Example 1. In this example, we simulate a channel with two states, 1 and 2. At each state, the availability periods are i.i.d. random variables exponentially distributed with $\mu_1 = 1/4$ and $\mu_2 = 1$. Also, the data unit sizes are continuous random variables, following the exponential distribution with mean 2 ($\lambda = 1/2$). Therefore, by definition, we have $\alpha_1 = 2$ and $\alpha_2 = 0.5$. The transition probability from state *i* to state *j* is defined as p_{ij} so that the steady state probabilities are given by $\pi_i = p_{ji}/(p_{ij} + p_{ji})$, i = 1, 2. In Fig. 3, we present the asymptotics of the number of retransmissions on the logarithmic scale for three values of steady state probabilities: $\pi_1 = 0.1, \pi_1 = 0.5$ and $\pi_1 = 0.8$; observe that m = 1



Figure 3: Example 1(b). Asymptotics of $\mathbb{P}(T > t)$ for a two-state channel.

is the index of the state with the larger α . We plot the exact asymptotics from Theorem 2.2, where we note that the constant term $\Gamma(\alpha_m + 1)/\pi_m^{m}$ increases the precision of our logarithmic asymptotics (Theorem 2.1). We observe that our simulation results are in excellent agreement with the theoretical asymptote.

Next, for the same channel and two values of π_1 , namely 0.1, and 0.5, Fig. 3 demonstrates the logarithmic asymptotics for $\mathbb{P}(T > t)$ obtained from Theorem 2.3.



Figure 4: Example 2. Asymptotics of $\mathbb{P}(N>n)$ for a three-state channel.

Example 2. In this example, we consider a three state channel, with transition probabilities such that $\pi_m = 0.25$ and $\pi_m = 0.75$, m = 1. The availability periods at each state are exponentially distributed with $\mu_1 = 2, \mu_2 = 1/2$ and $\mu_3 = 1/4$ and the data sizes are exponential with unit mean. From Fig. 4, we observe that the lightest asymptotics (power law with exponent $\alpha = 4$) dominates the tail of $\mathbb{P}(N > n)$. The power law tail appears relatively early when $\pi_m = 0.75$ and is not affected by other states for almost the entire



Figure 5: Example 3. Logarithmic Asymptotics for a two-state channel where data sizes and channel statistics are normally distributed.

Example 3. In the third example, we simulate a twostate channel where the packet sizes and the availability periods are normally distributed and the channel alternates between the two states with probability 1/2. Suppose that A and L take absolute values of zero mean normal random variables, with $\sigma_L = 5$ and $\sigma_{A_1} = 2, \sigma_{A_2} = 4, 6, 8$, for states 1 and 2, respectively. Note that $\mathbb{P}[|N(0, \sigma^2)| > x] \sim$ $2\sigma/(\sqrt{2\pi}x)e^{-x^2/2\sigma^2}$ and thus $\log \bar{G}(x) \approx \sigma_A^2/\sigma_L^2 \log \bar{F}(x)$, i.e. the asymptotic assumption of Theorem 2.1 is satisfied. In Fig. 5, we plot the logarithmic asymptotics for three different values of $\alpha = \sigma_A^2/\sigma_L^2$, as marked on the graph.



Figure 6: Example 4. Simulation of a two-state channel with bounded data sizes.

Example 4. In our last example, we simulate a two-state channel in the practically important case where the packet

sizes are bounded. In this setting, it is easy to see that the distribution of N has an exponential tail. However, this tail may appear for very small probabilities, which implies that the number of retransmissions (or delay) of interest can fall inside the main body of the distribution that behaves as power law. We simulate the channel of our first example, but here we further assume that L has finite support [0, b]; we consider two values for b, e.g. 5 and 10. In Fig. 6, we plot the distributions of $\mathbb{P}(N > n)$, parameterized by b, in two different scenarios where the stationary probability of the 'best' state π_m is 0.1, and 0.01 respectively. From the figure we observe that, when we increase the support of the distribution from b = 5 to b = 10, the region of the main (power law) body of the distribution of N increases, e.g. for $\pi_m = 0.1$ there is a tenfold increase, from less than 100 to almost 10^3 . This effect was first discussed in [9] and then rigorously proved in [6]. Nevertheless, the transition to the exponential behavior is not insensitive to the probabilities π_m . For the same b, higher stationary probabilities of the 'best' state tend to shrink the power law region. Hence, the exponential behavior happens at different times depending on how frequent the 'good' state is; we provide a deeper insight on this in the following Section 4.

4. ENGINEERING IMPLICATIONS AND CONCLUDING REMARKS

In this section, we discuss the engineering implications of our results. Previously, we showed that when the channel is correlated, meaning that it switches between dependent states, the 'best case' scenario wins. This implies that the delay asymptotics and the stability conditions are determined by the state that generates the lightest tail in the corresponding independent model. This insensitivity to the detailed structure of the correlations as well as the optimistic best case predictions are beneficial both for modeling and dimensioning/capacity planning of such systems. In particular, our insights show that the independent channel model might be sufficient. Furthermore, the analysis of the independent model is more likely to be extended to more complex multi-channel and networking systems with failures.

4.1 A Word of Caution

However, in this subsection, we emphasize that a design relying on the best case scenario may be overly optimistic and even completely wrong if the best state of the channel is atypical, i.e. it occurs very rarely. To illustrate this point, we study the following simplified model that demonstrates the impact of the correlated channel states on the distribution of N. In particular, we consider a channel with two states (K = 2), such that $\alpha_1 > \alpha_2$, and we assume that π_1 is very small ($\pi_1 \ll \pi_2, \pi_1 \approx 0$), i.e. state 1 is much less frequent than state 2. In this case, the tail of the distribution is still a power law with index $\alpha_m = \alpha_1$. However, there exists another power law asymptote that appears earlier and dominates the body of the distribution for smaller values of n.

We herein characterize these two asymptotes with two equivalent explicit formulas that approximate the retransmission distribution for large n (informal derivations are presented in the Appendix):

$$\mathbb{P}(N > n) \approx \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \sum_{i=0}^{\infty} \frac{(-n^{1-\delta} P_2)^i \Gamma(\alpha_2 + i\delta)}{i!}, \qquad (6)$$

and

$$\mathbb{P}(N>n) \approx \frac{\alpha_1}{n^{\alpha_1} \pi_1^{\alpha_1}} \sum_{i=0}^{\infty} \frac{(-n^{1-1/\delta} P_1)^i \Gamma(\alpha_1 + i/\delta)}{i!}, \quad (7)$$

where $\delta = \alpha_2/\alpha_1, P_1 = \pi_2/\pi_1^{1/\delta}$ and $P_2 = \pi_1/\pi_2^{\delta}$. Note that the sum in (6) is absolutely convergent since

Note that the sum in (6) is absolutely convergent since $\Gamma(\alpha_2 + i\delta) \leq \lceil \alpha_2 + i\delta \rceil!$. As we can easily infer from the first expression, when $n^{1-\delta}P_2 \ll 1$, the leading term dominates and the initial part of the distribution is determined by the heavier power law $O(n^{-\alpha_2})$ with exponent $\alpha_2 < \alpha_1$. This is indeed the asymptote that works well for small values of n, specifically when $n^{1-\delta} \ll 1/P_2$, as will be evident in the forthcoming example. Accordingly, the leading term of the second asymptote from (7) yields the correct tail asymptotics from Theorem 2.2, which holds as $n \to \infty$.

In order to illustrate the preceding derivation, we plot the exact asymptotes from equations (6) and (7), in Fig. 7, and compare with simulation. In this setting, we assume exponential distributions as in Example 1, such that $\alpha_1 = 2$ and $\alpha_2 = 1/2$; the steady state probabilities are $\pi_1 = 0.01$ and $\pi_2 = 0.99$. Specifically, we use five error terms for both asymptotes. We observe that the precision of the first asymptote deteriorates after $n \approx 10^2$ unless we increase the number of terms in expression (6). The leading term is a power law of index 1/2, which leads to the heavier asymptote. On the other hand, the tail asymptote derived in this section, even with few terms, fits perfectly for large values of n, which lends credit to our main Theorem 2.2.



Figure 7: Exact asymptotes from (6) and (7) for a two-state channel where $\alpha_1 = 2$ and $\alpha_2 = 1/2$.

In conclusion, our results imply that the tail distribution of the delay in a dynamic channel will be light as long as there is at least one state that generates light tail asymptotics. However, using the tail as a primary performance measure might result in an optimistic design that will expose the system to high congestion if the 'good' state is relatively rare. In such a case, the main body of the distribution may be characterized by different power laws and a mixture of distributions in between that are, in principle, much heavier than the tail. This situation must be treated with caution in order to guarantee system stability and good performance for all n. As illustrated in the last example, the main body can exhibit power law asymptotics with index $\alpha < 1$, which corresponds to a system with zero throughput. If the design does not account for this behavior, it is highly likely that the system will achieve poor performance for a considerably long period of time.

In general, the best strategy under unpredictable situations when the channel is unknown is to utilize channel feedback, possibly combined with dynamic fragmentation based on the number of unsuccessful retransmission attempts.

5. PROOFS

In this section, we present the proof for the upper bound of Theorem 2.1 as well as the proofs of Theorems 2.2, 2.3. **Proof:** [of Theorem 2.1]

Here, we prove the *upper bound*. Using the notation $\mathbb{E}_L[\cdot] = \mathbb{E}[\cdot|L]$ and similarly as before

$$\mathbb{P}(N > n|L) = \mathbb{E}_L \left[\prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{N_n^k} \right]$$

$$\leq \mathbb{E}_L \left[\prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{(1-\epsilon)\pi_k n} \mathbf{1} \left\{ N_n^k \ge (1-\epsilon)\pi_k n \right\} \right]$$

$$+ \mathbb{E}_L \left[\prod_{k=1}^K \mathbf{1} \left\{ N_n^k \le (1-\epsilon)\pi_k n \right\} \right]$$

$$= \prod_{k=1}^K \mathbb{P} \left(N_n^k \ge (1-\epsilon)\pi_k n \right) \mathbb{P}(L > A|J = k, L)^{(1-\epsilon)\pi_k n}$$

$$+ \prod_{k=1}^K \mathbb{P} \left(N_n^k \le (1-\epsilon)\pi_k n \right),$$

where $\mathbb{P}(N_n^k \ge (1-\epsilon)\pi_k n) \to 1$, as $n \to \infty$, implied by ergodicity and stationarity, whereas, from our main assumption, $\prod_{k=1}^{K} \mathbb{P}(N_n^k \le (1-\epsilon)\pi_k n) = O(1/n^{\alpha_m+\epsilon})$. Thus,

$$\mathbb{P}(N > n) = \mathbb{E}[\mathbb{P}(N > n|L)]$$

$$\leq (1+\epsilon)^{K} \mathbb{E} \prod_{k=1}^{K} \left(1 - \bar{F}(L)^{\frac{1}{\alpha_{k}(1-\epsilon)}}\right)^{\pi_{k}n(1-\epsilon)} \mathbf{1}\{L \ge x_{0}\}$$

$$+ (1+\epsilon)^{K} \prod_{k=1}^{K} \left(1 - \bar{F}(x_{0})^{\frac{1}{\alpha_{k}(1-\epsilon)}}\right)^{\pi_{k}n(1-\epsilon)} + \frac{\epsilon}{n^{\alpha_{m}+\epsilon}},$$

where x_0 is such so that (2) holds. Now using the elementary inequality $1 - x \leq e^x$, and by picking n_0 large so that $(1 + \epsilon)^K \left(1 - \bar{F}(x_0)^{\frac{1}{\alpha_k(1-\epsilon)}}\right)^{n(1-\epsilon)} \leq \epsilon/n^{\alpha_m+\epsilon}$, for $n \geq n_0$, we obtain

$$\mathbb{P}(N > n) \leq (1+\epsilon)^{K} \mathbb{E} \prod_{k=1}^{K} e^{-\pi_{k}n(1-\epsilon)\bar{F}(L)\frac{1}{\alpha_{k}(1-\epsilon)}} \mathbf{1}\{L \geq x_{0}\}$$
$$+ \frac{2\epsilon}{n^{\alpha_{m}+\epsilon}}$$
$$\leq (1+\epsilon)^{K} \mathbb{E} \left[e^{-\sum_{k=1}^{K} \pi_{k}n(1-\epsilon)U\frac{1}{\alpha_{k}(1-\epsilon)}} \right] \mathbf{1}\{L \geq x_{0}\}$$
$$+ \frac{\epsilon}{n^{\alpha_{m}(1-\epsilon)}},$$

where the first integral is derived similarly as in the proof of the lower bound, by picking x_0 large. Thus

$$\mathbb{P}(N > n) \le \frac{(1-\epsilon)\Gamma(\alpha_m(1-\epsilon)+1)}{(\pi_m n)^{\alpha_m(1-\epsilon)}} + \frac{\epsilon}{n^{\alpha_m(1-\epsilon)}}.$$

Next, we set $H_{\epsilon} = (1 - \epsilon)\Gamma(\alpha_m(1 - \epsilon) + 1)/\pi_m^{\alpha_m(1 - \epsilon)} + \epsilon$, and after taking the logarithm, we obtain

 $\log \mathbb{P}(N > n) \le \log H_{\epsilon} - \alpha_m (1 - \epsilon) \log n,$

and by picking n_0 such that $\log H_{\epsilon} \leq \alpha_m \epsilon \log n$, we have

$$\log \mathbb{P}(N > n) \le -\alpha_m (1 - 2\epsilon) \log n.$$

Last, replacing ϵ with $\epsilon/2$ yields

$$\frac{\log \mathbb{P}(N > n)}{\log n} \le -(1 - \epsilon)\alpha_m. \tag{8}$$

- - - - 1 - - 1

Letting $\epsilon \to 0$ in both (5) and (8) completes the proof. \Box **Proof:** [of Theorem 2.2]

We begin with the *lower bound*. Following the same arguments as in the proof of Theorem 2.1, we obtain

. .

$$\mathbb{P}(N > n|L) = \mathbb{P}[L > A_1, L > A_2 \dots, L > A_n|L]$$

$$= \mathbb{E}_L[\mathbb{P}(L > A_j, 1 \le j \le n|J_1, \dots, J_n, L)]$$

$$= \mathbb{E}_L\left[\prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{N_n^k}\right]$$

$$\geq \mathbb{E}_L\left[\prod_{k=1}^K \mathbb{P}(L > A|J = k, L)^{(1+\epsilon)\pi_k n} \mathbf{1}\left\{N_n^k \le (1+\epsilon)\pi_k n\right\}\right]$$

$$\geq (1-\epsilon)^K \prod_{k=1}^K (1-\bar{G}_k(L))^{(1+\epsilon)\pi_k n},$$

which follows from recalling that $\mathbb{P}(N_n^k \leq (1 + \epsilon)\pi_k n) \geq (1 - \epsilon)$ for $n \geq n_0$.

Next, note that by assumption $\Phi_k(\cdot)$ is monotonic since $\overline{F}(x)^{-1} = \Phi_k(\overline{G}_k^{-1}(x))$ and both $\overline{F}(x), \overline{G}_k(x)$ are monotonic. Hence, we can define (in a standard way) an inverse $\Phi_k^{\leftarrow}(\cdot)$ such that $\Phi_k^{\leftarrow}(\Phi_k(x)) = x$, where $\Phi_k^{\leftarrow}(x)$ is also regularly varying (e.g. see Theorem 1.5.11 in [2]). Now, by Proposition 1.5.8 of [2], we can always find an absolutely continuous, strictly monotone and locally bounded (for x > 0) function

$$\Phi_k^*(x) = \begin{cases} \alpha \int_1^x \Phi_k(s) s^{-1} ds, & x \ge 1\\ 0, & 0 \le x < 1, \end{cases}$$
(9)

which satisfies $\Phi_k(x) \sim \Phi_k^*(x)$ for x large.

Next, by Theorem 1.5.13 in [2], $\Phi_k^{\leftarrow}(x)$ is asymptotically equivalent to the inverse of $\Phi_k^*(x)$; observe that $\Phi_k^*(x)$ is strictly monotone, for all $x \geq 1$, which guarantees that its inverse exists. Hence, our main assumption becomes

$$\bar{G}_k(x)^{-1} \sim \Phi_k^{*\leftarrow}(\bar{F}(x)^{-1}),$$

where $\Phi_k^{*\leftarrow}(\cdot)$ is absolutely continuous (differentiable) for each k. This implies that, for any $0 < \epsilon < 1$ and $x \ge x_0$, we have

$$(1-\epsilon)/\Phi_k^{*\leftarrow}\left(\bar{F}(x)^{-1}\right) \le \bar{G}_k(x) \le (1+\epsilon)/\Phi_k^{*\leftarrow}\left(\bar{F}(x)^{-1}\right),$$
(10)

In the rest of the proof, for simplicity, we will abuse the notation and simply write $\Phi_k(\cdot)$ and $\Phi_k^{\leftarrow}(\cdot)$ to denote $\Phi_k^*(\cdot)$ and its inverse $\Phi_k^{*\leftarrow}(\cdot)$, respectively.

Therefore,

$$\mathbb{P}(N > n) = \mathbb{E}[P(N > n|L)]$$

$$\geq (1 - \epsilon)^{K} \mathbb{E} \prod_{k=1}^{K} \left(1 - \frac{1 + \epsilon}{\Phi_{k}^{\leftarrow}(\bar{F}(L)^{-1})}\right)^{\pi_{k}n(1+\epsilon)} \mathbf{1}\{L \ge x_{n}\}$$

where, we can choose x_n such that $\Phi_k^{\leftarrow}(\bar{F}(x_n)^{-1}) = n/H$, for *n* large and H > 0. Thus, by the elementary inequality $1 - x \ge e^{-(1+\epsilon)x}$ for small *x*, we obtain

$$\begin{split} \mathbb{P}(N > n) \\ &\geq (1 - \epsilon)^{K} \mathbb{E} \prod_{k=1}^{K} \exp\left(-\frac{n\pi_{k}(1 + \epsilon)^{3}}{\Phi_{k}^{\leftarrow}(U^{-1})}\right) \mathbf{1}\{U \leq \bar{F}(x_{n})\} \\ &= (1 - \epsilon)^{K} \int_{0}^{\bar{F}(x_{n})} \exp\left(-n\sum_{k=1}^{K} \frac{\pi_{k}(1 + \epsilon)^{3}}{\Phi_{k}^{\leftarrow}(u^{-1})}\right) du \\ &\geq (1 - \epsilon)^{K} \int_{0}^{\bar{F}(x_{n})} \exp\left(-\frac{n(1 + \epsilon)^{4}\pi_{m}}{\Phi_{m}^{\leftarrow}(u^{-1})}\right) du, \end{split}$$

where we observe that $\overline{F}(L) = U$, where U is uniform in (0,1), and that for large n, $\sum_{k \neq m} \pi_k \Phi_m^{\leftarrow}(u^{-1})/\pi_m \Phi_k^{\leftarrow}(u^{-1}) \leq \epsilon$, for small u. Next, by changing the variables $z = n/\Phi_m^{\leftarrow}(u^{-1})$, we obtain for small h > 0,

$$\mathbb{P}(N>n) \ge (1-\epsilon)^K \int_h^H e^{-(1+\epsilon)^4 \pi_m z} \frac{\Phi'_m(n/z)}{\Phi_m^2(n/z)} \frac{n}{z^2} dz$$

Now, we use the properties of regularly varying functions that, for positive h, H, we obtain uniformly for all $h \leq z \leq H$ and large n,

$$\frac{\Phi_m(n)}{\Phi_m(n/z)} \ge (1-\epsilon)z^{\alpha_m}$$
 and $\frac{\Phi_m'(n/z)}{\Phi_m(n/z)} = \frac{\alpha_m z}{n},$

where the equation follows from (9), for n > h. Thus,

$$\mathbb{P}[N > n] \ge \frac{(1-\epsilon)^{K+1}\alpha_m}{\Phi_m(n)} \int_h^H e^{-(1+\epsilon)^4 \pi_m z} z^{\alpha_m - 1} dz$$
$$\ge \frac{(1-\epsilon)\alpha_m}{\pi_m^{\alpha_m} \Phi_m(n)} \int_h^{H\pi_m} e^{-y} y^{\alpha_m - 1} dy,$$

which is derived after replacing $(1-\epsilon)^{K+1}/(1+\epsilon)^{4\alpha_m}$ with $(1-\epsilon)$ and by change of variables. Last, letting $H \to \infty$ and $h \to 0$, we obtain

$$\mathbb{P}(N > n) \ge (1 - \epsilon) \frac{\Gamma(\alpha_m + 1)}{\pi_m^{\alpha_m} \Phi_m(n)},\tag{11}$$

for $n \ge n_0$, which proves the lower bound.

For the *upper bound*, similarly as before, we obtain

$$\mathbb{P}(N > n|L) \leq \mathbb{E}\left[\prod_{k=1}^{K} (1 - \bar{G}_k(L))^{(1-\epsilon)\pi_k n} \mathbf{1}\left\{N_n^k \geq (1-\epsilon)\pi_k n\right\}\right] \\ + \mathbb{E}\left[\prod_{k=1}^{K} \mathbf{1}\left\{N_n^k \leq (1-\epsilon)\pi_k n\right\}\right] \\ \leq (1+\epsilon)^K \prod_{k=1}^{K} (1 - \bar{G}_k(L))^{(1-\epsilon)\pi_k n} + \prod_{k=1}^{K} \mathbb{P}\left(N_n^k \leq (1-\epsilon)\pi_k n\right),$$

where, by ergodicity and stationarity, we recall that as $n \to \infty$, $\mathbb{P}(N_n^k \ge (1-\epsilon)\pi_k n) \to 1$ whereas, from our main as-

sumption,
$$\prod_{k=1}^{K} \mathbb{P}(N_n^k \ge (1-\epsilon)\pi_k n) = O(1/n^{\alpha_m+\epsilon}).$$
 Thus,
$$\mathbb{P}(N > n) = \mathbb{E}[\mathbb{P}(N > n|L)]$$
$$\le (1+\epsilon)^K \mathbb{E} \prod_{k=1}^{K} \left(1 - \frac{1-\epsilon}{\Phi_k^{\leftarrow}(\bar{F}(L)^{-1})}\right)^{\pi_k n(1-\epsilon)} \mathbf{1}\{L \ge x_0\}$$
$$+ (1+\epsilon)^K \mathbb{E} \prod_{k=1}^{K} \left(1 - \bar{G}_k(x_0)\right)^{\pi_k n(1-\epsilon)} + \frac{\epsilon}{n^{\alpha_m+\epsilon}}$$
$$:= I_1 + \frac{2\epsilon}{n^{\alpha_m+\epsilon}},$$

where we pick x_0 such that $(1 + \epsilon)^K (1 - \bar{G}(x_0))^{\pi_k n(1-\epsilon)} \le \epsilon/n^{\alpha_m+\epsilon}$, for $n \ge n_0$. Now using the elementary inequality $1 - x \le e^x$, we have

$$I_{1} \leq (1+\epsilon)^{K} \mathbb{E} \exp\left(-\sum_{k=1}^{K} \frac{\pi_{k} n(1-\epsilon)^{2}}{\Phi_{k}^{\leftarrow}(\bar{F}(L)^{-1})}\right)$$
$$\leq (1+\epsilon)^{K} \int_{0}^{1} \exp\left(-\sum_{k=1}^{K} \frac{\pi_{k} n(1-\epsilon)^{2}}{\Phi_{k}^{\leftarrow}(1/u)}\right) du$$
$$\leq (1+\epsilon)^{K} \int_{0}^{1} \exp\left(-\frac{\pi_{m} n(1-\epsilon)^{2}}{\Phi_{m}^{\leftarrow}(1/u)}\right) du,$$

where we argue similarly as in the preceding proof for the lower bound. Then, changing the variables $z = n/\Phi_m^{\leftarrow}(1/u)$ yields

$$I_{1} \leq (1+\epsilon)^{K} \int_{0}^{1} e^{-\frac{\pi_{m} n(1-\epsilon)^{2}}{\Phi_{m}^{\leftarrow}(1/u)}} \mathbf{1} \left(h \leq \frac{n}{\Phi_{m}^{\leftarrow}(u^{-1})} \leq e^{m}\right) du + \sum_{k=m}^{\lceil \log(n/n\epsilon) \rceil} e^{-\pi_{m} e^{k}} \mathbb{P} \left[e^{k} \leq \frac{n}{\Phi_{m}^{\leftarrow}(U^{-1})} \leq e^{k+1}\right] + e^{-n/n\epsilon} := I_{11} + I_{12} + I_{10}.$$

First, we compute I_{11} as

$$I_{11} \leq (1+\epsilon)^K \int_h^{e^m} e^{-\pi_m (1-\epsilon)^2 z} \frac{\Phi'_m(n/z)n}{\Phi_m^2(n/z)z^2} dz$$
$$\leq \frac{(1+\epsilon)^{K+1}\alpha_m}{\Phi_m(n)} \int_h^{e^m} e^{-\pi_m (1-\epsilon)^2 z} z^{\alpha_m - 1} dz$$
$$\leq \frac{(1+\epsilon)\alpha_m}{\pi_m^{\alpha_m} \Phi_m(n)} \int_{h(1-\epsilon)^2 \pi_m}^{e^m} e^{-z} z^{\alpha_m - 1} dz,$$

where we replace $(1 + \epsilon)^{K+1}/(1 - \epsilon)^{2\alpha_m}$ with $(1 + \epsilon)$. Now, I_{12} becomes

$$I_{12} \leq \sum_{k=m}^{\lceil \log(n/n_{\epsilon}) \rceil} \frac{e^{-\pi_m e^k}}{\Phi_m(n/e^{k+1})}$$
$$\leq \sum_{k=m}^{\lceil \log(n/n_{\epsilon}) \rceil} \frac{e^{-\pi_m e^k}(1+\epsilon)^{k+1}}{\Phi_m(n)} \leq \frac{\epsilon}{\Phi_m(n)},$$

since the preceding sum is finite and $\Phi_m(n)/\Phi_m(n/e^k) \leq (1+\epsilon)^k$ for all $n \geq n_0$.

Last, I_{10} decays exponentially fast and thus $I_{10} = o(1/\Phi_m(n))$. Therefore,

$$\mathbb{P}(N>n)\Phi_m(n)\pi_m^{\alpha_m} \le (1+\epsilon)\alpha_m \int_0^{e^m} e^{-z} z^{\alpha_m-1} dz + o(1).$$
(12)

Performance Evaluation Review, Vol. 41, No. 2, September 2013

Note that passing $m \to \infty$ in the first term of (12), yields that for all $n > n_0$,

$$\frac{\mathbb{P}(N>n)\Phi_m(n)\pi_m^{\alpha_m}}{\Gamma(\alpha_m+1)} \le 1+2\epsilon.$$

After replacing 2ϵ with $\epsilon/2$, we obtain the upper bound

$$\mathbb{P}(N > n) \le (1 + \epsilon) \frac{\Gamma(\alpha_m + 1)}{\pi_m^{\alpha_m} \Phi_m(n)},\tag{13}$$

which, along with (11), finishes the proof. \Box **Proof:** [of Theorem 2.3]

In the following proof, we use the notation $(x \land y) = \min(x, y)$ to refer to the minimum of x and y. First we prove the *upper* bound.

For any $0 < \delta < 1$, we have

$$\begin{split} \mathbb{P}[T > t] &= \mathbb{P}\left[\sum_{i=1}^{N-1} A_i + L > t\right] \\ &\leq \mathbb{P}\left[\sum_{i=1}^{N-1} (A_i \wedge L) > (t - t^{1-\delta}), L \leq t^{1-\delta}\right] + \mathbb{P}[L > t^{1-\delta}] \\ &\leq \mathbb{P}\left[\sum_{i=1}^{N} (A_i \wedge t^{1-\delta}) > (t - t^{1-\delta})\right] + \mathbb{P}[L > t^{1-\delta}] \\ &\leq \mathbb{P}\left[\sum_{i=1}^{N} (A_i \wedge t^{1-\delta}) > (1 - \epsilon)t, N \leq t^{1-\delta}\right] \\ &+ \mathbb{P}\left[N > t^{1-\delta}\right] + \mathbb{P}[L > t^{1-\delta}] \\ &:= I_1 + I_2 + I_3, \end{split}$$

where in the third inequality, we use $t - t^{1-\delta} \approx (1-\epsilon)t$ for large t. First, I_3 is upper bounded by

$$I_3 \le \frac{\mathbb{E}[L^{\alpha_m} \mathbf{1}\{L > t^{1-\delta}\}]}{t^{\alpha_m(1-\delta)}} = o(1/t^{\alpha_m(1-\delta)}), \qquad (14)$$

since the condition $\mathbb{E}(A^{1+\theta}) < \infty$, together with our main assumption, implies that for x_0 as in (2),

$$\mathbb{E}[L^{\alpha_m}] = x_0 + \int_{x_0}^{\infty} \mathbb{P}[L^{\alpha_m} > x] dx \le x_0$$
$$+ \int_{x_0}^{\infty} \mathbb{P}[A > x^{1/\alpha_m} | J = m]^{\alpha_m(1-\epsilon)} dx$$
$$\le x_0 + \int_{x_0}^{\infty} \frac{(\mathbb{E}A^{1+\theta})^{\alpha_m(1-\epsilon)}}{x^{(1+\theta)(1-\delta)(1-\epsilon)}} dx < \infty,$$

as $t \to \infty$ for $1 + \theta > 1/(1 - \delta)(1 - \epsilon)$. Next, for I_1 , we have

$$I_{1} = \mathbb{EP}\left[\sum_{i=1}^{t^{1-\delta}} (A_{i} \wedge t^{1-\delta}) > (1-\epsilon)t|J_{1}, J_{2}, \dots, J_{\lfloor t^{1-\delta} \rfloor}\right]$$

$$\leq \mathbb{E}e^{-\theta(1-\epsilon)t}\mathbb{E}\exp\left[\sum_{i=1}^{t^{1-\delta}} \theta(A_{i} \wedge t^{1-\delta}) > (1-\epsilon)t|J_{1}, \dots, J_{\lfloor t^{1-\delta} \rfloor}\right]$$

$$= \mathbb{E}e^{-\theta(1-\epsilon)t}\prod_{i=1}^{t^{1-\delta}}\mathbb{E}\exp\left[\theta(A_{i} \wedge t^{1-\delta})|J_{i}\right],$$

which follows by applying the exponential Chebyshev's inequality for $\theta > 0$. Now, observe that $\mathbb{E}[A_i|J_i] \leq$ $\max_{k=1,\dots,K} \mathbb{E}[A|J = k] =: \mu_m \text{ and using the inequality}$

Performance Evaluation Review, Vol. 41, No. 2, September 2013

 $e^x \leq 1 + xe^y, 0 \leq x \leq y$, we upper bound the exponential moments of $X_i := (A_i \wedge t^{1-\delta} | J_i)$ by

$$\mathbb{E}e^{\theta(A_i \wedge t^{1-\delta})|J_i} \le 1 + e^{\theta t^{1-\delta}} \theta \mathbb{E}(A_i|J_i)$$
$$\le 1 + e^{\theta t^{1-\delta}} \theta \mu_m \le \exp\left(\theta \mu_m e^{\theta t^{\delta-1}}\right),$$

which renders

$$I_{1} \leq e^{-\theta(1-\epsilon)t} \exp\left(t^{1-\delta}\theta\mu_{m}e^{\theta t^{\delta-1}}\right)$$
$$= e^{-(1-\epsilon)t^{\delta}}e^{\mu_{m}e} \leq o\left(t^{-\alpha_{m}(1-\delta)}\right), \qquad (15)$$

where we pick $\theta = t^{\delta - 1}$.

From Theorem 2.2, we recall that for $0 < \delta < 1$,

$$\lim_{t \to \infty} \frac{\log \mathbb{P}\left[N > t^{1-\delta}\right]}{\log t} = -(1-\delta)\alpha_m,$$

which, along with (14) and (15), and passing $\delta \to 0$, implies

$$\limsup_{t \to \infty} \frac{\log \mathbb{P}[T > t]}{\log t} \le -\alpha_m.$$
(16)

Next, we establish the *lower bound*. It follows easily that

$$\mathbb{P}[T > t] = \mathbb{P}\left[\sum_{i=1}^{N-1} A_i + L > t\right]$$

$$\geq \mathbb{P}\left[\sum_{i=1}^{N-1} A_i > t, N \ge t^{1+\delta} + 1\right]$$

$$\geq \mathbb{P}\left[N \ge t^{1+\delta} + 1\right] - \mathbb{P}\left[\sum_{i=1}^{t^{1+\delta}} A_i \le t\right]$$

$$\geq \mathbb{P}\left[N \ge t^{1+\delta} + 1\right] - \mathbb{E}\left[\mathbb{P}\left[\sum_{i=1}^{t^{1+\delta}} A_i \le t | J_1, J_2, \dots, J_{\lfloor t^{1-\delta} \rfloor}\right]\right]$$

$$:= I_1 - I_2.$$

Now, we can show that $I_2 \leq o(t^{-\alpha_m(1+\delta)})$, by similar arguments as in the proof of (15) for the upper bound; we omit the details.

Regarding I_1 , we recall from Theorem 2.1 that for 0 < 1 $\delta < 1$, we have

$$\lim_{t \to \infty} \frac{\log \mathbb{P}\left[N > t^{1+\delta} + 1\right]}{\log t} = -(1+\delta)\alpha_m,$$

and thus, by passing $\delta \to 0$,

$$\liminf_{t \to \infty} \frac{\log \mathbb{P}[T > t]}{\log t} \ge -\alpha_m.$$
(17)

Finally, combining (16) and (17) concludes the proof. \Box

- $\begin{array}{c} \textbf{6,} \quad \textbf{REFERENCES} \\ [1] \quad \mathrm{S. Asmussen, P. Fiorini, L. Lipsky, T. Rolski, and} \end{array}$ R. Sheahan. Asymptotic behavior of total times for jobs that must start over if a failure occurs. Mathematics of Operations Research, 33(4):932-944, November 2008.
- [2] N. H. Bingham, C. M. Goldie, and J. L. Teugels. Regular Variation. Cambridge University Press, 1987.
- [3] P. M. Fiorini, R. Sheahan, and L. Lipsky. On unreliable computing systems when heavy-tails appear as a result of the recovery procedure. SIGMETRICS Performance Evaluation Review, 33(2):15–17, December 2005.
- [4] P. R. Jelenkovic and A. A. Lazar. Subexponential asymptotics of a markov-modulated random walk with queueing applications. Technical report, 1996.

- [5] P. R. Jelenković and E. D. Skiani. Distribution of the number of retransmissions of bounded documents. submitted to Advances in Applied Probability, New York, 2012, eprint arXiv:1210.8421v1.
- [6] P. R. Jelenković and E. D. Skiani. Uniform approximation of the distribution for the number of retransmissions of bounded documents. In *Proceedings of the 12th ACM SIGMETRICS/PERFORMANCE joint international conference on Measurement and Modeling of Computer Systems*, SIGMETRICS '12, pages 101–112, June 2012.
- [7] P. R. Jelenković and J. Tan. Stability of finite population ALOHA with variable packets. *Technical Report EE2009-02-20, Department of Electrical Engineering, Columbia University*, New York, 2009, eprint arxiv: 0902.4481v2.
- [8] P. R. Jelenković and J. Tan. Are end-to-end acknowlegements causing power law delays in large multi-hop networks? In 14th Informs Applied Probability Conference, July 2007.
- P. R. Jelenković and J. Tan. Can retransmissions of superexponential documents cause subexponential delays? In *Proceedings of IEEE INFOCOM'07*, pages 892–900, May 2007.
- [10] P. R. Jelenković and J. Tan. Is ALOHA causing power law delays? In Proceedings of the 20th International Teletraffic Congress, Ottawa, Canada ; Lecture Notes in Computer Science, No 4516, Springer-Verlag, pages 1149–1160, June 2007.
- [11] P. R. Jelenković and J. Tan. Dynamic packet fragmentation for wireless channels with failures. In *Proceedings of the 9th* ACM International Symposium on Mobile Ad Hoc Networking and Computing, May 2008.
- [12] P. R. Jelenković and J. Tan. Modulated branching processes, origins of power laws, and queueing duality. *Mathematics of Operations Research*, 35(4):807–829, 2010.
- [13] P. R. Jelenković and J. Tan. Characterizing heavy-tailed distributions induced by retransmissions. Advances in Applied Probability, 45(1): 106-138, 2013. (extended version) arXiv: 0709.1138v2.
- [14] J. Nair, M. Andreasson, L. Andrew, S. Low, and J. Doyle. File fragmentation over an unreliable channel. In *Proceedings of IEEE INFOCOM'10*, pages 965–973, March 2010.
- [15] S. M. Ross. A First Course in Probability. Prentice Hall, 6th edition, 2002.
- [16] R. Sheahan, L. Lipsky, P. M. Fiorini, and S. Asmussen. On the completion time distribution for tasks that must restart from the beginning if a failure occurs. *SIGMETRICS Performance Evaluation Review*, 34(3):24–26, December 2006.

APPENDIX

Appendix

In this Appendix, we provide the informal derivation of formulas (6) and (7) of Subsection 4.1.

Proof: [of (6) and (7)]

Starting from equation (3), assuming that $\bar{F}(x) \approx \bar{G}_k(x)^{\alpha_k}$, k = 1, 2, and using similar arguments as in the derivation of (3)-(4), we informally argue that

$$\begin{split} \mathbb{P}(N > n) &\approx \mathbb{E}[(1 - \bar{G}_1(L))^{\pi_1 n} (1 - \bar{G}_2(L))^{\pi_2 n}] \\ &\approx \mathbb{E}[(1 - \bar{F}(L)^{1/\alpha_1})^{\pi_1 n} (1 - \bar{F}(L)^{1/\alpha_2})^{\pi_2 n}] \\ &\approx \mathbb{E}[(1 - U^{1/\alpha_1})^{\pi_1 n} (1 - U^{1/\alpha_2})^{\pi_2 n}] \\ &\approx \mathbb{E}[e^{-\pi_1 n U^{1/\alpha_1} - \pi_2 n U^{1/\alpha_2}}], \end{split}$$

since $\overline{F}(L) = U$, where U is uniformly distributed in (0,1). Thus,

$$\mathbb{P}(N > n) \approx \int_{0}^{1} e^{-\pi_{1}nu^{1/\alpha_{1}} - \pi_{2}nu^{1/\alpha_{2}}} du$$

$$= \frac{\alpha_{2}}{n^{\alpha_{2}}\pi_{2}^{\alpha_{2}}} \int_{0}^{n\pi_{2}} e^{-z - \frac{\pi_{1}nz^{\alpha_{2}/\alpha_{1}}}{(n\pi_{2})^{\alpha_{2}/\alpha_{1}}}} z^{\alpha_{2}-1} dz,$$
(18)

which follows by changing the variables $z = \pi_2 n u^{1/\alpha_2}$. Now, let $\delta := \alpha_2/\alpha_1 < 1$ and $P_2 := \pi_1/\pi_2^{\delta}$, so that

$$\begin{split} \mathbb{P}(N>n) &\approx \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \int_0^{n\pi_2} e^{-z - n^{1-\delta} P_2 z^{\delta}} z^{\alpha_2 - 1} dz \\ &= \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \int_0^{n\pi_2} e^{-z} z^{\alpha_2 - 1} \left(1 - n^{1-\delta} P_2 z^{\delta} \right. \\ &\quad + \frac{(n^{1-\delta} P_2)^2 z^{2\delta}}{2} - \dots + \frac{(-n^{1-\delta} P_2)^i z^{i\delta}}{i!} + \dots \right) dz, \end{split}$$

by the Taylor expansion of the function $e^x = \sum_{i=0}^{\infty} x^i/i!$. Now by extending the integral to infinity we have

$$\begin{split} \mathbb{P}(N > n) &\approx \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \int_0^\infty e^{-z} z^{\alpha_2 - 1} \left(1 - n^{1 - \delta} P_2 z^{\delta} \right. \\ &+ \frac{(n^{1 - \delta} P_2)^2 z^{2\delta}}{2} - \dots + \frac{(-n^{1 - \delta} P_2)^i z^{i\delta}}{i!} + \dots \right) dz \\ &= \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \left(\int_0^\infty e^{-z} z^{\alpha_2 - 1} dz - n^{1 - \delta} P_2 \int_0^\infty e^{-z} z^{\alpha_2 + \delta - 1} dz \right. \\ &+ \frac{(n^{1 - \delta} P_2)^2}{2} \int_0^\infty e^{-z} z^{\alpha_2 + 2\delta - 1} dz - \\ &\cdots + \frac{(-n^{1 - \delta} P_2)^i}{i!} \int_0^\infty e^{-z} z^{\alpha_2 + i\delta - 1} dz + \dots \right) \\ &= \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \left(\Gamma(\alpha_2) - n^{1 - \delta} P_2 \Gamma(\alpha_2 + \delta) \right. \\ &+ \frac{(n^{1 - \delta} P_2)^2 \Gamma(\alpha_2 + 2\delta)}{2} - \dots + \frac{(-n^{1 - \delta} P_2)^i \Gamma(\alpha_2 + i\delta)}{i!} + \dots \right), \end{split}$$

which follows immediately from the definition of the gamma function $\Gamma(\alpha) = \int_0^\infty e^{-z} z^{\alpha-1} dz$. This yields the explicit form

$$\mathbb{P}(N > n) \approx \frac{\alpha_2}{n^{\alpha_2} \pi_2^{\alpha_2}} \sum_{i=0}^{\infty} \frac{(-n^{1-\delta} P_2)^i \Gamma(\alpha_2 + i\delta)}{i!}.$$

By the same approach, starting from (18) and changing of variables as $z = n\pi_1 u^{\alpha_1}$, we obtain

$$\mathbb{P}(N > n) \approx \int_0^1 e^{-\pi_1 n u^{1/\alpha_1} - \pi_2 n u^{1/\alpha_2}} du$$
$$= \frac{\alpha_1}{n^{\alpha_1} \pi_1^{\alpha_1}} \int_0^{n\pi_1} e^{-z - \frac{\pi_2 n z^{\alpha_1/\alpha_2}}{(n\pi_1)^{\alpha_1/\alpha_2}}} z^{\alpha_1 - 1} dz$$

Now, for $P_1 := \pi_2/\pi_1^{1/\delta}$ and using Taylor expansion of e^x , we have

$$\begin{split} \mathbb{P}(N>n) &\approx \frac{\alpha_1}{n^{\alpha_1} \pi_1^{\alpha_1}} \int_0^{n\pi_1} e^{-z - n^{1-1/\delta} P_1 z^{1/\delta}} z^{\alpha_1 - 1} dz \\ &\approx \frac{\alpha_1}{n^{\alpha_1} \pi_1^{\alpha_1}} \int_0^{\infty} e^{-z} z^{\alpha_1 - 1} \left(1 - n^{1-1/\delta} P_1 z^{1/\delta} \right. \\ &\left. + \frac{(n^{1-1/\delta} P_1)^2 z^{2/\delta}}{2} - \dots + \frac{(-n^{1-1/\delta} P_1)^i z^{i/\delta}}{i!} + \dots \right) dz, \end{split}$$

for large n. By identical arguments as before,

$$\mathbb{P}(N>n) \approx \frac{\alpha_1}{n^{\alpha_1} \pi_1^{\alpha_1}} \left(\Gamma(\alpha_1) - n^{1-1/\delta} P_1 \Gamma(\alpha_1 + 1/\delta) + \frac{(n^{1-1/\delta} P_1)^2 \Gamma(\alpha_1 + 2/\delta)}{2} - \dots + \frac{(-n^{1-1/\delta} P_1)^i \Gamma(\alpha_1 + i/\delta)}{i!} + \dots \right),$$

which yields the explicit form

$$\mathbb{P}(N>n)\approx \frac{\alpha_1}{n^{\alpha_1}\pi_1^{\alpha_1}}\sum_{i=0}^\infty \frac{(-n^{1-1/\delta}P_1)^i\Gamma(\alpha_1+i/\delta)}{i!}. \quad \Box$$