Can Retransmissions of Superexponential Documents Cause Subexponential Delays?

Predrag R. Jelenković Jian Tan
Department of Electrical Engineering
Columbia University, New York, NY 10027
{predrag, jiantan}@ee.columbia.edu
August 2006

***Abstract*— Consider a generic data unit of random size L that needs to be transmitted over a channel of unit capacity. The channel dynamics is modeled as an on-off process $\{(A_i, U_i)\}_{i\geq 1}$ with alternating independent periods when channel is available A_i and unavailable U_i, respectively. During each period of time that the channel becomes available, say A_i, we attempt to transmit the data unit. If $L \leq A_i$, the transmission was considered successful; otherwise, we wait for the next period A_{i+1} when the channel is available and attempt to retransmit the data from the beginning. We study the asymptotic properties of the total transmission time T and number of retransmissions N until the data is successfully transmitted.**

In recent studies [1], [2], it was proved that the waiting time T follows a power law when the distributions of L and A_1 are of an exponential type, e.g., Gamma distribution. In this paper, we show that the distributions of N and T follow power laws with exponent α as long as $\log \mathbb{P}[L > x] \approx \alpha \log \mathbb{P}[A_1 > x]$ for large x. Hence, it may appear surprising that we obtain power law distributions irrespective of how heavy or light the distributions of L and A_1 may be. In particular, both L and A_1 can decay faster than any exponential, which we term superexponential. For example, if L and A_1 are Gaussian with variances σ_L^2 and σ_A^2, respectively, then N and T have power law distributions with exponent $\alpha = \sigma_A^2/\sigma_L^2$; note that, if $\sigma_A^2 < \sigma_L^2$, the transmission time has an infinite mean and, thus, the system is unstable.

The preceding model, as recognized in [1], describes a variety of situations where failures require jobs to restart from the beginning. Here, we identify that this model also provides a new mechanism for explaining the frequently observed power law phenomenon in data networks. Specifically, we argue that it may imply the power laws on both the application as well as the data link layer, where variable-sized documents and (IP) packets are transmitted, respectively. We discuss the engineering ramifications of our observations, especially in the context of wireless ad hoc and sensor networks where channel failures are frequent. Furthermore, our results provide an easily computable benchmark for measuring the matching between the data and channel characteristics that permits/prevents satisfactory transmission.

I. Introduction

In this paper, we study a problem of transmitting a generic data unit of random size L over a channel with failures. For example, most of us have experienced difficulties with downloading files/Web pages from the Internet where in the middle of the download the connection breaks and one has to restart the download from the beginning. Similar problems might arise in transmitting (IP) packets of variable length over the data link layer. Here, in case of an error, the packet is resent automatically, which is known as the Automatic Repeat reQuest (ARQ) protocol (e.g., see Section 2.4 of [3]). Both of the preceding problems are especially important in the wireless environment, e.g., for wireless ad hoc and sensor networks, where frequent channel failures occur due to a variety of reasons, including interference, signal fading, contention with other nodes, multipath effects, obstructions, node mobility, and other changes in the environment [4].

We use the following generic channel with failures to model the preceding situations. The channel dynamics is described as an on-off process $\{(A_i, U_i)\}_{i\geq 1}$ with alternating independent periods when channel is available A_i and unavailable U_i, respectively. In each period of time that the channel becomes available, say A_i, we attempt to transmit the data unit of random size L. If $L \leq A_i$, we say that the transmission was successful; otherwise, we wait for the next period A_{i+1} when the channel is available and attempt to retransmit the data from the beginning. We study the asymptotic properties of the distributions of the total transmission time T and number of retransmissions N, for the precise definitions of these variables and the model, see Section II.

The preceding model was introduced and studied in [5] and, apart from the already mentioned applications in communications, it represents a generic model for other situations where jobs have to restart from the beginning after a failure. It was first recognized in [1] that this model results in power laws when the distributions of L and A_1 have a matrix exponential representation. The more recent study in [2] rigorously proves that the distribution of T is asymptotically a power law when the variables L and A_1 are of exponential type (e.g., Gamma distribution).

In this paper, we discover that the distributions of N and T follow power laws with the same exponent α as long as $\log \mathbb{P}[L > x] \approx \alpha \log \mathbb{P}[A_1 > x]$ for large x. The recognition that this relative condition between the distributions of L and A_1 is the primary cause of power laws, rather than any absolute assumptions (e.g., exponential), represents one of the main novelties of our approach. Hence, maybe somewhat surprisingly, one obtains power law distributions irrespective

This work was supported by the National Science Foundation under grant number CNS-0435168.

Published in Proceedings of INFOCOM'2007, Anchorage, Alaska, May 6-12, 2007.

of how heavy or light the distributions of L and A_1 are, allowing both L and A_1 to decay faster than any exponential, which we term superexponential; see Theorems 2 and 3 of Section III-A. In particular, if L and A_1 take absolute values of Gaussian random variables with variances σ_L^2 and σ_A^2, respectively, then N and T have power law distributions with exponent $\alpha = \sigma_A^2/\sigma_L^2$; note that, if $\sigma_A^2 < \sigma_L^2$, the transmission time has an infinite mean and, thus, the system is unstable. Furthermore, in Section III-B, we refine our results by establishing the exact power law asymptotics in Theorems 4 and 5. Finally, in Section III-C, we show that, if the distribution of L has an infinite support ($\mathbb{P}[L > x] > 0$ for all $x \geq 0$), the distributions of N and T are subexponential (decay slower than any exponential) regardless of what the distribution of A_1 may be.

Power laws, and in general heavy tails, are frequently observed in computer communication networks and systems. The heavy-tailed transmission delays are typically attributed to power law distributions of files/documents on the internet, e.g., see [6] and the references therein. Hence, transmitting these heavy-tailed objects over communication channels (that have bounded capacity) naturally results in power law delays. In this paper, we recognize that the power laws can result entirely from the retransmissions, a commonly used component of communication protocols, even if the documents as well as the channel characteristics are light-tailed (i.e., bounded by an exponential).

In Section IV, we illustrate our theoretical results with simulation and numerical experiments. In particular, we emphasize the characteristics of the studied channel that may not be immediately apparent from our theorems. For example, the relative logarithmic condition that we identify as a cause of power laws is based on high order distributional properties and, thus, it is quite insensitive to the mean values of L and A_1. Interestingly enough, we show that, even if the expected data size $\mathbb{E}L$ is much smaller than the average length of channel availability $\mathbb{E}A_1$, the transmission delays can be power laws with infinite expected delays and number of retransmissions. Furthermore, in practice, the distribution of documents/packets might have a bounded support. We show that this situation may result in distributions of T and N exhibiting power laws in the main body, i.e., essentially truncated power law distributions. To this end, it is also important to note that the power law main body have an exponentiated (stretched) support in relation to the support of L and, thus, may result in very long, although, exponentially bounded delays.

From an engineering perspective, our main discovery is the matching between the statistical characteristics of the channel and transmitted data (packets). Basically, if $\log \mathbb{P}[A > x] > \log \mathbb{P}[L > x]$ or $\log \mathbb{P}[A > x] < \log \mathbb{P}[L > x]$, then one can expect good or bad (infinite mean) delay performance. We discuss these and other engineering implications of our results in Section V. In particular, we focus on the wireless environment where channel/connection failures are frequent. As stated earlier, most of us have been inconvenienced when the connections would brake while we are downloading a large file from the Internet. This issue has been already recognized in practice where software for downloading was developed that would save the intermediate data (checkpoints) and resume the download from the point when the connection was broken. However, our results emphasize that, in the presence of frequently failing connections, the long delays may arise even when downloading relatively small documents. Hence, we argue that one may need to adopt the application layer software for the wireless environment by introducing checkpoints even for small to moderate size documents. We also discuss possible large delays that may result on the data link layer due to (IP) packet variability and channel failures. We assume that in the physical layer, the codewords, which represent the basic units of packet transmission, are much smaller than the maximum size of the packet. We believe that this is a realistic situation for sensor networks, where complicated coding schemes are unlikely since the nodes have very limited computational power. In this context, our results show that the number of retransmissions could be power law, which challenges the traditional model that assumes a geometric number of retransmissions. We discuss possible solutions to alleviate this problem, such as breaking large packets into smaller units. Obviously there is a tradeoff between the sizes of these newly created packets and the throughput since, if the packets are too small, they will mostly contain the packet headers and, thus, very little useful information.

II. Description of the Channel

In this section, we formally describe our model and provide necessary definitions and notation.

Let L denote the random length of a generic data unit (packet). Without loss of generality, we assume that the channel is of unit capacity. The channel dynamics is modeled as an on-off process $\{(A_i, U_i)\}_{i\geq 1}$ with alternating independent periods when channel is available A_i and unavailable U_i, respectively. In each period of time that the channel becomes available, say A_i, we attempt to transmit the data unit and, if $L \leq A_i$, we say that the transmission was successful; otherwise, we wait for 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.

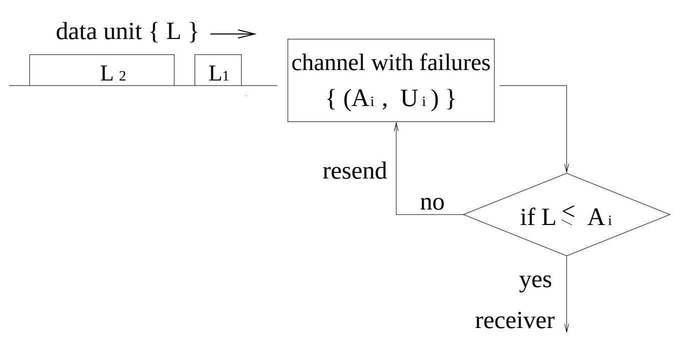

Fig. 1. Packets sent over channels with failures

Assume that $\{U, U_i\}_{i\geq 1}$ and $\{A, A_i\}_{i\geq 1}$ are two mutually independent sequences of i.i.d. random variables.

Definition 1: The total number of (re)transmissions for a generic data unit of length L is defined as

$$N \triangleq \inf\{n : A_n \geq L\},$$

and, the total transmission time for the data unit is defined as

$$T \triangleq \sum_{i=1}^{N-1}(A_i + U_i) + L.$$

We use the following notations to denote the complementary cumulative distribution functions for A and L respectively,

$$\bar{G}(x) \triangleq \mathbb{P}[A > x],$$

and

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

As already stated in the introduction, a problem of this type (with $U_i \equiv 0$), in a different application context, was defined and examined in [5]. It was first recognized in [1] that this model can result in power laws when the distributions of L and A have a matrix exponential representation. The more recent study in [2] rigorously proves that the distribution of T is asymptotically a power law when the variables L and A are of exponential type (e.g., Gamma distribution). The following theorem is quoted from Theorem 7 in [2]. In this paper we use the following standard notations. For any two real functions $a(t)$ and $b(t)$ and fixed $t_0 \in \mathbb{R} \cup \{\infty\}$, 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 $\underline{\lim}_{t\to t_0} a(t)/b(t) \geq 1$; $a(t) \lesssim b(t)$ has a complementary definition.

Theorem 1 (AFLS06): Assume that $\bar{G}(x) = e^{-\beta x}$ and $\bar{F}(x) \sim (a/\delta)x^b e^{-\delta x}$ where $b \in \mathbb{R}$ and $a, \beta, \delta > 0$, $U_i \equiv 0$, then,

$$\mathbb{P}[T > t] \sim a\Gamma(\alpha)\,\beta^{-(\alpha+b+1)}\frac{(\log t)^b}{t^\alpha}, \tag{1}$$

where $\alpha = \delta/\beta$.

In the following section we will derive more general results of this type. The preceding theorem will be a direct consequence of our Theorem 5. The main novelty of our results is that it reveals that power law arises from the relative value between the hazard functions of L and A, and does not depend on the absolute forms of G and F. Hence, G and F can have arbitrarily heavy or light tails, as long as, roughly speaking, their hazard functions are asymptotically proportional, e.g., see equation (2) in the forthcoming Theorem 2.

III. Main Results

This section presents our main results. Here, we assume that $\bar{F}(x)$ is a continuous function with support on $[0, \infty)$. The same results can be derived when $\bar{F}(x)$ is lattice valued which we will prove in the extended version of this paper [7]. However, some of the examples in Section IV will be based on discrete valued random variables.

A. Logarithmic Power Law Asymptotics

In this subsection we present the logarithmic asymptotics for the number of transmissions and the total transmission time in Theorem 2 and Theorem 3, respectively. These results imply that logarithmic scale is the right measure of the interplay between the data unit size and the lengths of channel availability. The proof of Theorem 3 is deferred to Section VI.

Theorem 2: If there exists $\alpha > 0$, such that,

$$\lim_{x\to\infty}\frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \alpha, \tag{2}$$

then, we have

$$\lim_{n\to\infty}\frac{\log \mathbb{P}[N > n]}{\log n} = -\alpha. \tag{3}$$

Theorem 3: Under the same condition of Theorem 2 and $\mathbb{E}[(U + A)^{1+\alpha+\theta}] < \infty$ for some $\theta > 0$, then,

$$\lim_{n\to\infty}\frac{\log \mathbb{P}[T > t]}{\log t} = -\alpha. \tag{4}$$

Remark 1: These two theorems indicate that the distribution tails of the number of transmissions and total transmission time are essentially power laws. Thus, the system can exhibit high variations and possible instability, e.g., when $1 < \alpha < 2$, the transmission time has an infinite variance and, when $0 < \alpha < 1$, it does not even have a finite mean.

Remark 2: It is easy to understand that if the data sizes (e.g., files, packets) follow heavy-tailed distributions, the total transmission time is also heavy-tailed. However, from these two theorems, we see that even if the distributions of the data and channel characteristics are highly concentrated, i.e., light-tailed (e.g., see Corollary 1 below), once they are asymptotically proportional on the logarithmic scale, the heavy-tailed transmission delays can still arise.

1) Proof of Theorem 2: Notice that the number of retransmissions is geometrically distributed given the packet size L,

$$\mathbb{P}[N > n \mid L] = (1 - \bar{G}(L))^n,$$

and, therefore,

$$\mathbb{P}[N > n] = \mathbb{E}[(1 - \bar{G}(L))^n]. \tag{5}$$

First, let us establish an upper bound. The condition described in (2) implies that for any $0 < \epsilon < 1/\alpha$, there exists x_ϵ, such that for all $x > x_\epsilon$, we have $\bar{F}(x)^{\frac{1}{\alpha}+\epsilon} \leq \bar{G}(x) \leq \bar{F}(x)^{\frac{1}{\alpha}-\epsilon}$. Hence,

$$\begin{aligned}\mathbb{E}[(1 - \bar{G}(L))^n] &= \mathbb{E}[(1 - \bar{G}(L))^n \mathbf{1}(L > x_\epsilon)] \\ &\quad + \mathbb{E}[\left(1 - \bar{G}(L)\right)^n \mathbf{1}(L \leq x_\epsilon)] \\ &\leq \mathbb{E}\left[\left(1 - \bar{F}(L)^{\frac{1}{\alpha}+\epsilon}\right)^n\right] + \left(1 - \bar{G}(x_\epsilon)\right)^n \\ &\leq \mathbb{E}\left[e^{-n\bar{F}(L)^{\frac{1}{\alpha}+\epsilon}}\right] + \left(1 - \bar{G}(x_\epsilon)\right)^n.\end{aligned}$$

Since $\bar{F}(x)$ is continuous, $\bar{F}(L)$ is a uniform random variable (denoted by U) between 0 and 1 (e.g., see Proposition 2.1 in

Chapter 10 of [8]), we derive

$$\mathbb{P}[N>n] \leq \mathbb{E}\left[e^{-nU^{\frac{1}{\alpha}+\epsilon}}\right] + \left(1-\bar{G}(x_\epsilon)\right)^n$$
$$= \frac{\Gamma\left(\alpha/(1+\alpha\epsilon)+1\right)}{n^{\alpha/(1+\alpha\epsilon)}} + \left(1-\bar{G}(x_\epsilon)\right)^n, \quad (6)$$

where the last equality is due to the identity

$$\mathbb{E}\left[e^{-\theta U^{1/\beta}}\right] = \Gamma(\beta+1)/\theta^\beta, \text{ for } \theta, \beta > 0. \quad (7)$$

Therefore, by letting $n \to \infty$, and then $\epsilon \to 0$ in (6), we obtain

$$\overline{\lim_{n\to\infty}} \frac{\log \mathbb{P}[N>n]}{\log n} \leq -\alpha. \quad (8)$$

Next, we derive a lower bound. Noticing that for any $0 < \delta < 1$, there exists $x_\delta > 0$ such that $1-x \geq e^{(1-\delta)x}$ for all $0 < x < x_\delta$, we can choose x_ϵ large enough, such that

$$\mathbb{E}[(1-\bar{G}(L))^n] \geq \mathbb{E}\left[\left(1-\bar{G}(L)\right)^n \mathbf{1}(L > x_\epsilon)\right]$$
$$\geq \mathbb{E}\left[\left(1-\bar{F}(L)^{\frac{1}{\alpha}-\epsilon}\right)^n \mathbf{1}\left(\bar{F}(L) < \bar{F}(x_\epsilon)\right)\right]$$
$$\geq \mathbb{E}\left[e^{-n(1-\epsilon)\bar{F}(L)^{\frac{1}{\alpha}-\epsilon}} \mathbf{1}\left(\bar{F}(L) < \bar{F}(x_\epsilon)\right)\right]$$
$$\geq \mathbb{E}\left[e^{-n(1-\epsilon)U^{\frac{1}{\alpha}-\epsilon}}\right] - e^{-n(1-\epsilon)\bar{F}(x_\epsilon)^{\frac{1}{\alpha}-\epsilon}}.$$

Recalling the identity (7), passing $n \to \infty$, and then $\epsilon \to 0$, we derive

$$\underline{\lim_{n\to\infty}} \frac{\log \mathbb{P}[N>n]}{\log n} \geq -\alpha. \quad (9)$$

Finally, combining (8) and (9), we obtain (3). ■

B. Exact Power Law Asymptotics

This section presents the exact asymptotic results for the number of retransmissions and the total transmission time. The proof of Theorem 5 is deferred to Section VI.

Theorem 4: If $\bar{G}(x) \sim H\left(-\log \bar{F}(x)\right) \bar{F}(x)^{1/\alpha}$, $\alpha > 0$ with $H(x)$ being a continuous and regularly varying function, then, as $n \to \infty$,

$$\mathbb{P}[N>n] \sim \frac{\Gamma(\alpha+1)}{n^\alpha H(\alpha \log n)^\alpha}. \quad (10)$$

Theorem 5: Under the same conditions as in Theorem 3 and 4, we have, as $t \to \infty$,

$$\mathbb{P}[T>t] \sim \frac{\Gamma(\alpha+1)(\mathbb{E}[U+A])^\alpha}{t^\alpha H(\alpha \log t)^\alpha}. \quad (11)$$

The preceding theorems, under a bit more restrictive conditions, characterize the exact asymptotic tail behavior of the distributions of N and T and, therefore, refine Theorems 2 and 3.

Note that Theorem 1 can be easily derived from Theorem 5 using the following argument. First, it is easy to verify that, as $x \to \infty$,

$$\bar{G}(x) \sim \frac{\delta^{(b+1)/\alpha}}{a^{1/\alpha}} \left(-\log \bar{F}(x)\right)^{-b/\alpha} \bar{F}(x)^{1/\alpha},$$

where $\alpha = \delta/\beta$, and, therefore, we can choose

$$H(x) = \frac{\delta^{(b+1)/\alpha}}{a^{1/\alpha}} x^{-b/\alpha}.$$

Thus, using Theorem 5, we derive the asymptotics in (1).

Before moving to the proof, we state one more straightforward consequence of the preceding theorems that allows $\bar{F}$ and $\bar{G}$ to have normal-like distributions, i.e., much lighter tails than exponential.

Corollary 1: Suppose $\bar{G}(x) = \mathbb{P}[|N(0,\sigma_A^2)| > x]$ and $\bar{F}(x) = \mathbb{P}[|N(0,\sigma_L^2)| > x]$, where $N(0,\sigma^2)$ is a Gaussian random variable with mean zero and variance σ^2, then,

$$\mathbb{P}[N>n] \sim \Gamma(\alpha+1)\, \alpha^{-1/2} \frac{(\pi \log n)^{\frac{1}{2}(\alpha-1)}}{n^\alpha}, \quad (12)$$

where $\alpha = {\sigma_A}^2/{\sigma_L}^2$.

Proof: First, notice that

$$\mathbb{P}[|N(0,\sigma^2)| > x] \sim \frac{2\sigma}{\sqrt{2\pi}x} e^{-\frac{x^2}{2\sigma^2}},$$

and, therefore, recalling that $\alpha = {\sigma_A}^2/{\sigma_L}^2$, it is easy to obtain

$$\bar{G}(x) \sim \pi^{\frac{1}{2}(1/\alpha-1)} \alpha^{1/2} \left(-\log \bar{F}(x)\right)^{\frac{1}{2}(1/\alpha-1)} \left(\bar{F}(x)\right)^{1/\alpha}.$$

Hence, $\bar{F}(x)$ and $\bar{G}(x)$ satisfy the assumption of Theorem 4 with

$$H(x) = \pi^{\frac{1}{2}(1/\alpha-1)} \alpha^{1/2} x^{\frac{1}{2}(1/\alpha-1)},$$

which implies (12). ■

1) Proof of Theorem 4: First, let us prove the upper bound. Observe that

$$\mathbb{P}[N>n] = \mathbb{E}\left[\left(1-\bar{G}(L)\right)^n\right]$$
$$\leq \mathbb{E}\left[e^{-n\bar{G}(L)}\right]$$
$$\leq \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(L < x_1)\right]$$
$$+ \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(x_1 \leq L \leq x_2)\right]$$
$$+ \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(L > x_2)\right]$$
$$\triangleq I_1 + I_2 + I_3. \quad (13)$$

Then, for any $1 > \delta > 0$, there exists x_δ, such that for all $x > x_\delta$, we have $(1-\delta)H\left(-\log \bar{F}(x)\right)\bar{F}(x)^{1/\alpha} \leq \bar{G}(x) \leq (1+\delta)H\left(-\log \bar{F}(x)\right)\bar{F}(x)^{1/\alpha}$. Next, since $\bar{F}(x)$ is continuous, there exist $x_1, x_2 > x_\delta$, such that $\bar{F}(x_2) = 1/n^{\alpha+\epsilon}$ and $\bar{F}(x_1) = 1/n^{\alpha-\epsilon}$ for any $\alpha > \epsilon > 0$. Therefore, as $n \to \infty$,

$$I_1 = \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(L < x_1)\right]$$
$$\leq e^{-n\bar{G}(x_1)}$$
$$\leq e^{-n(1-\delta)H\left(-\log \bar{F}(x_1)\right)\bar{F}(x_1)^{1/\alpha}}$$
$$= o\left(\frac{1}{n^\alpha}\right), \quad (14)$$

and,

$$I_3 = \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(L > x_2)\right]$$
$$\leq \mathbb{P}[L > x_2]$$
$$= o\left(\frac{1}{n^\alpha}\right). \quad (15)$$

Now, since $\bar{F}(x)$ and $H(x)$ are continuous, $\bar{F}(L) \triangleq U$ is a uniform random variable between 0 and 1 (e.g., see Proposition 2.1 in Chapter 10 of [8]), and there exists $\alpha - \epsilon \leq \xi_1 \leq \alpha + \epsilon$, such that $H(\xi_1 \log n) \leq H(-\log \bar{F}(L))$ when $x_1 \leq L \leq x_2$. Therefore,

$$\begin{aligned} I_2 &= \mathbb{E}\left[e^{-n\bar{G}(L)}\mathbf{1}(x_1 < L < x_2)\right] \\ &\leq \mathbb{E}\left[e^{-n(1-\delta)H\left(-\log \bar{F}(L)\right)\bar{F}(L)^{1/\alpha}}\mathbf{1}\left(x_1 < L < x_2\right)\right] \\ &\leq \mathbb{E}\left[e^{-n(1-\delta)H(\xi_1 \log n)U^{1/\alpha}}\right] \\ &= \frac{\Gamma(\alpha+1)}{(n(1-\delta)H\left(\alpha \log n\right))^{\alpha}}\frac{\left(H\left(\alpha \log n\right)\right)^{\alpha}}{\left(H\left(\xi_1 \log n\right)\right)^{\alpha}}, \end{aligned} \tag{16}$$

and, by the Characterisation Theorem of regular variation (e.g., see Theorem 1.4.1 of [9]) and the uniform convergence theorem of slowly varying functions (Theorem 1.2.1 of [9]), we have

$$\lim_{n\to\infty} \frac{H\left(\alpha \log n\right)}{H\left(\xi_1 \log n\right)} \to 1 \ \text{ as } \epsilon \to 0. \tag{17}$$

Hence, using (13), (14), (15), (16), (17) and passing δ, ϵ to zero, we derive

$$\mathbb{P}[N > n] \lesssim \frac{\Gamma(\alpha+1)}{n^{\alpha}H\left(\alpha \log n\right)^{\alpha}}. \tag{18}$$

Next, we prove the lower bound. For any $\epsilon > 0$, choose $x_1, x_2 > x_\delta$, such that $\bar{F}(x_2) = (\log n)^{-\epsilon}/n^{\alpha}$ and $\bar{F}(x_1) = (\log n)^{\epsilon}/n^{\alpha}$. Since $H(x)$ is continuous, there exists $\alpha - \epsilon \log\log n/\log n \leq \xi_2 \leq \alpha + \epsilon \log\log n/\log n$, such that $H(\xi_2 \log n) \geq H(-\log \bar{F}(L))$ when $x_1 \leq L \leq x_2$. Therefore,

$$\begin{aligned} \mathbb{P}[N > n] &= \mathbb{E}\left[\left(1 - \bar{G}(L)\right)^n\right] \\ &\geq \mathbb{E}\left[\left(1 - \bar{G}(L)\right)^n \mathbf{1}(x_1 \leq L \leq x_2)\right] \\ &\geq \mathbb{E}\Big[\left(1 - (1+\delta)H\left(-\log \bar{F}(L)\right)\bar{F}(L)^{1/\alpha}\right)^n \\ &\qquad \mathbf{1}(x_1 \leq L \leq x_2)\Big] \\ &\geq \mathbb{E}\Bigg[\left(1 - (1+\delta)H\left(\xi_2 \log n\right)U^{1/\alpha}\right)^n \\ &\qquad \mathbf{1}\left(\frac{(\log n)^{-\epsilon}}{n^{\alpha}} \leq U \leq \frac{(\log n)^{\epsilon}}{n^{\alpha}}\right)\Bigg] \\ &= \int_{(\log n)^{-\epsilon/\alpha}}^{(\log n)^{\epsilon/\alpha}} \left(1 - \frac{(1+\delta)H\left(\xi_2 \log n\right)u}{n}\right)^n \frac{d\left(u^{\alpha}\right)}{n^{\alpha}} \\ &\sim \frac{\Gamma(\alpha+1)}{(1+\delta)^{\alpha}n^{\alpha}H\left(\xi_2 \log n\right)^{\alpha}} \ \text{ as } n \to \infty. \end{aligned} \tag{19}$$

Thus, by passing δ, ϵ to zero in (19), we derive

$$\mathbb{P}[N > n] \gtrsim \frac{\Gamma(\alpha+1)}{n^{\alpha}H\left(\alpha \log n\right)^{\alpha}}. \tag{20}$$

Combining (18) and (20) completes the proof. ∎

C. Subexponential asymptotics

The preceding results establish the relationships between L and A, under which power law tails arise. However, it is natural to ask if there is any other relationship between L and A, such that the distributions of N and T are light-tailed, i.e., bounded by an exponential. Interestingly enough, the following lemma gives a negative answer to this question. This result was proven in Theorem 6 of [2] under the assumption that $\bar{G}(x)$ is exponential, but, as it can be seen from Lemma 1, the exponential assumption is not necessary.

Lemma 1: If $\bar{F}(x) > 0$ for all $x \geq 0$, then both N and T are *subexponential* in the following sense that, for any $\epsilon > 0$,

$$e^{\epsilon n}\mathbb{P}[N > n] \to \infty \ \text{ as } n \to \infty, \tag{21}$$

and

$$e^{\epsilon t}\mathbb{P}[T > t] \to \infty \ \text{ as } t \to \infty. \tag{22}$$

Proof: Note that for any $1 > \delta > 0$, there exists $t_\delta > 0$ such that

$$1 - t \geq (1-\delta)e^{-t},$$

for all $0 < t < t_\delta$. Therefore, we can choose x_δ large enough, such that $1 - \bar{G}(x) \geq (1-\delta)e^{-\bar{G}(x)}$, for all $x > x_\delta$. Then,

$$\begin{aligned} e^{\epsilon n}\mathbb{P}[N > n] &\geq e^{\epsilon n}\mathbb{E}\left[\left(1 - \bar{G}(L)\right)^n \mathbf{1}(L \geq x_\delta)\right] \\ &\geq e^{\epsilon n}\mathbb{E}\left[(1-\delta)^n e^{-n\bar{G}(L)}\mathbf{1}(L \geq x_\delta)\right] \\ &\geq \left(e^{\epsilon - \bar{G}(x_\delta)}(1-\delta)\right)^n \bar{F}(x_\delta). \end{aligned}$$

Thus, by selecting δ small enough and x_δ large enough, we can always make $e^{\epsilon - \bar{G}(x_\delta)}(1-\delta) > 1$, and, by passing $n \to \infty$, we complete the proof of (21).

Next, suppose that $\bar{G}(x_0) > 0$ for some $x_0 > 0$; otherwise, T will be infinite, which yields (22) immediately. We can always find $x_1 > x_0 > 0$, such that i.i.d. random variables $X_i \triangleq x_0\mathbf{1}(x_0 < A_i < x_1)$ satisfy $0 < \mathbb{E}X_1 < \infty$. Now, for any $\zeta > 0$,

$$\begin{aligned} \mathbb{P}[T > t] &= \mathbb{P}\left[\sum_{i=1}^{N-1}(U_i + A_i) + L > t\right] \\ &\geq \mathbb{P}\left[\sum_{i=1}^{N-1} A_i\mathbf{1}(x_0 < A_i < x_1) > t\right] \\ &\geq \mathbb{P}\left[\sum_{i=1}^{N-1} X_i > t\right] \\ &\geq \mathbb{P}\left[\sum_{i=1}^{N-1} X_i > t, N \geq \frac{t(1+\zeta)}{\mathbb{E}X_1}\right] \\ &\geq \mathbb{P}\left[N > \frac{t(1+\zeta)}{\mathbb{E}X_1} + 1\right] \\ &\quad - \mathbb{P}\left[\sum_{i=1}^{N-1} X_i \leq t, N > \frac{t(1+\zeta)}{\mathbb{E}X_1} + 1\right] \\ &\triangleq I_1 - I_2. \end{aligned} \tag{23}$$

Since, for $\bar{X}_i \triangleq \mathbb{E}[X_i] - X_i$,

$$I_2 \le \mathbb{P}\left[\sum_{i \le t(1+\zeta)/\mathbb{E}X_1} X_i \le t\right] = \mathbb{P}\left[\sum_{i \le t(1+\zeta)/\mathbb{E}X_1} \bar{X}_i \ge \zeta t\right], \quad (24)$$

it is well known (e.g., see Example 1.15 of [10]) that there exists $\eta > 0$, such that

$$I_2 \le e^{-\eta t}. \quad (25)$$

Therefore, by (21), (23) and (25), we obtain that for all $0 < \epsilon < \eta$,

$$e^{\epsilon t}\mathbb{P}[T > t] \to \infty \text{ as } t \to \infty,$$

implying that (22) holds for any $\epsilon > 0$. ■

Hence, in view of our preceding results, it remains to characterize the situations when the distributions of N and T decay faster than polynomial but slower than exponential; these types of distributions are often referred to as being moderately heavy. However, since these distributions have all the moments finite, they can not be a cause of instability, and therefore, we defer this analysis to the extended version of the paper [7].

IV. Numerical and Simulation Examples

In this section, we illustrate our theoretical results with simulation and numerical experiments. In particular, we emphasize the characteristics of the studied channel that may not be immediately apparent from our theorems. For example, the relative logarithmic condition that we identify as a cause of power laws is based on higher order distributional properties and, thus, it is quite insensitive to the mean values of L and A. Interestingly enough, we show that, even if the expected data size $\mathbb{E}L$ is much smaller than the average length of channel availability $\mathbb{E}A$, the transmission delays can be power laws with infinite expected delays and retransmissions. Furthermore, in practice, the distribution of documents/packets might have a bounded support. We show that this situation may result in truncated power law distributions for T and N. To this end, it is also important to note that the distributions of N and T will have a power law main body with a stretched support in relation to the support of L and, thus, may result in very long, although, exponentially bounded delays.

Example 1: This example illustrates the exact asymptotic results presented in Theorem 4. We choose two sets of distributions. One assumes that A and L take absolute values of zero mean normal random variables, as stated in Corollary 1, with $\sigma_A = 4$ and $\sigma_L = 6$, respectively; the other assumes that A and L follow exponential distributions with parameters $\lambda_A = 6$ and $\lambda_L = 4$, respectively. Now, by Theorem 4 (Corollary 1), the asymptotic behavior of $\mathbb{P}[N > n]$ is given by $\Gamma(13/9)(3/2)(\pi \log n)^{-5/18} n^{-4/9}$, for the normal case; and by $\Gamma(5/3)n^{-2/3}$, for the exponential case. The simulation results for $50,000$ samples and the corresponding asymptotes are plotted on log-log scale in Figure 2. From this figure, we see that even for small values of n, say $n \approx 5$ for exponential distributions and $n \approx 10$ for normal distributions, the numerical asymptote approximates the simulation quite well and for larger values of n the simulation results and the asymptotic formulas are basically indistinguishable.

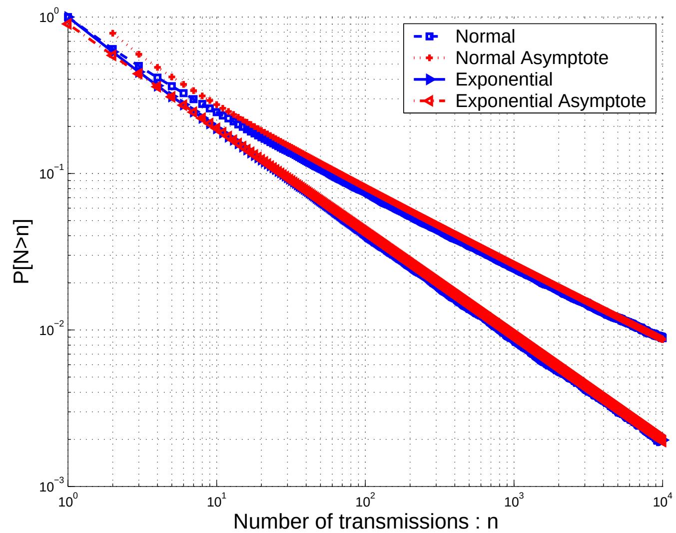

Fig. 2. First illustration for Example 1.

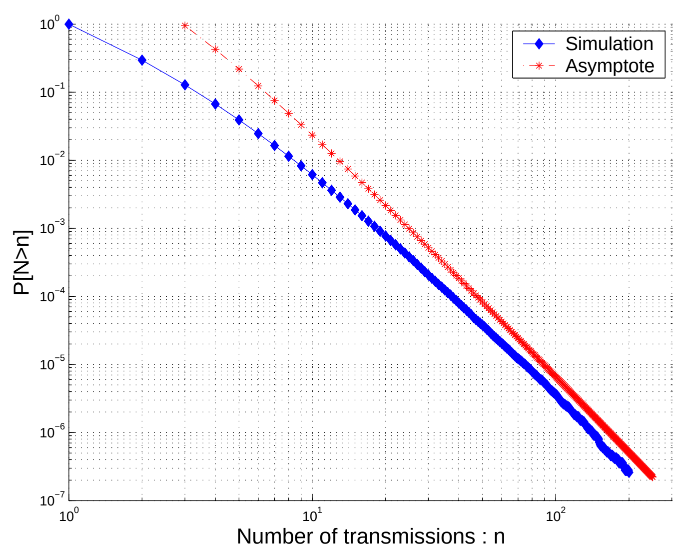

Fig. 3. Second illustration for Example 1.

In general, the asymptotic approximations may not be as accurate for small n (large probabilities) as in the previous cases. Here, we demonstrate this point with A and L taking absolute values of zero mean normal random variables with $\sigma_A = 8$ and $\sigma_L = 4$, respectively. The simulation results of 10^7 samples and the asymptote are plotted in Figure 3. From the figure we can see that only for $n > 100$ $\left(\mathbb{P}[N > n] < 10^{-5}\right)$ the asymptotic approximation becomes accurate. However, even in this situation the derived asymptote provides a right order of magnitude and shows the very heavy (highly variable) behavior of the distribution of N.

Example 2: Intuitively, one would expect that, if the expected data size $\mathbb{E}L$ is smaller than the average length of channel availability $\mathbb{E}A$, the system should behave reasonably well. In this regard, surprisingly, this example shows that not only that the distributions of N and T can have very heavy tails, but the system can even be unstable $\mathbb{E}N = \infty$, $\mathbb{E}T = \infty$. Suppose that $\bar{G}(x) = \mathbb{P}[(N(10, 6^2))^+ > x]$ and $\bar{F}(x) = \mathbb{P}[(N(15, 4^2))^+ > x]$, where $N(\mu, \sigma^2)$ is a Gaussian random variable with mean μ, variance σ^2 and x^+ denotes the positive part of x. Obviously $\mathbb{E}[L] \approx 10 < \mathbb{E}[A] \approx 15$, but we still get a power law distribution with $\alpha \approx 0.67 < 1$ for the number of transmissions N, which implies that N has an infinite mean. The simulation result for 5×10^5 samples is presented in Figure 4; the matching asymptote is drawn on the same figure with a dashed line.

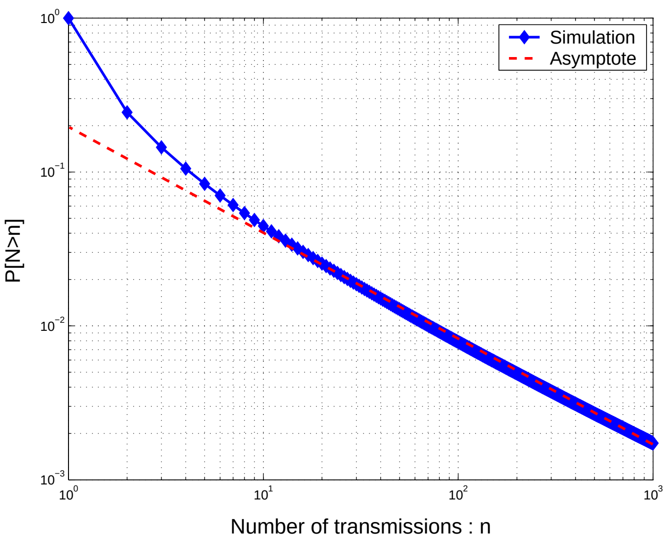

Fig. 4. Illustration for Example 2.

Example 3: In many practical applications the data unit L can be bounded, i.e., the distribution of F has a bounded support. Thus, from equation (5) it is easy to see that the distribution of N is exponentially bounded. However, this exponential behavior can happen for very small probabilities, while the number of retransmissions and delay of interest can fall inside the region of the distribution (main body) that behaves as power laws. This example is aimed to explain this important phenomenon. We assume that L has finite support $[0, K]$ and show how this results in a truncated power law distribution for N in the main body, even though the tail is exponentially bounded. This example is parametrized by K where K ranges from 20 to 30 and we choose the same parameters as in Example 2 except that whenever L exceed K, we set L to K. We plot the distributions of $\mathbb{P}[N > n]$, parameterized by K, in Figure 1. From the figure we can see that, when we increase the support of the distribution from $K = 20$ to $K = 30$, the main (power law) body of the

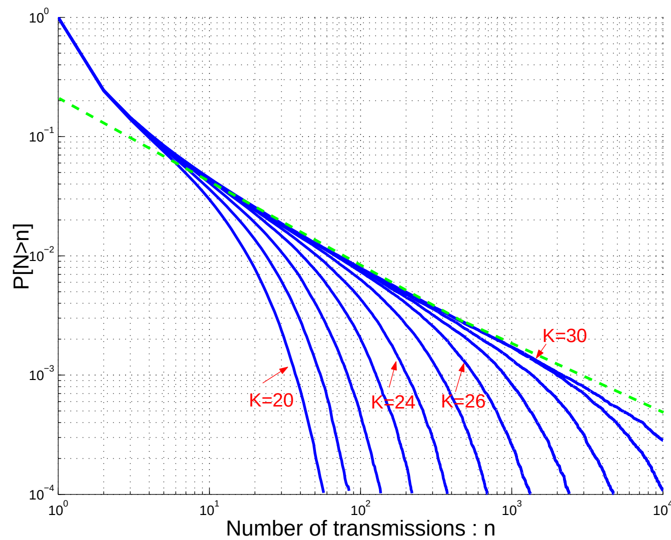

Fig. 5. Illustration for Example 3.

distribution of N increase from less than 10 to almost 10^4. This effect is what we call the stretched support of the main body of $\mathbb{P}[N > n]$ in relation to the support K of L. In fact, it can be rigorously shown that the support of the main body of $\mathbb{P}[N > n]$ grows quicker than an exponential function if the distributions of L and A are lighter than exponential. This is why we also refer to the support of the main body of $\mathbb{P}[N > n]$ as being exponentiated in relation to K. We will present this result formally in the extended version of this paper [7]. Furthermore, it is important to note that, if $K = 30$ and the probabilities of interest for $\mathbb{P}[N > n]$ are greater than 10^{-3}, then the results of this experiments are essentially indistinguishable from Example 2.

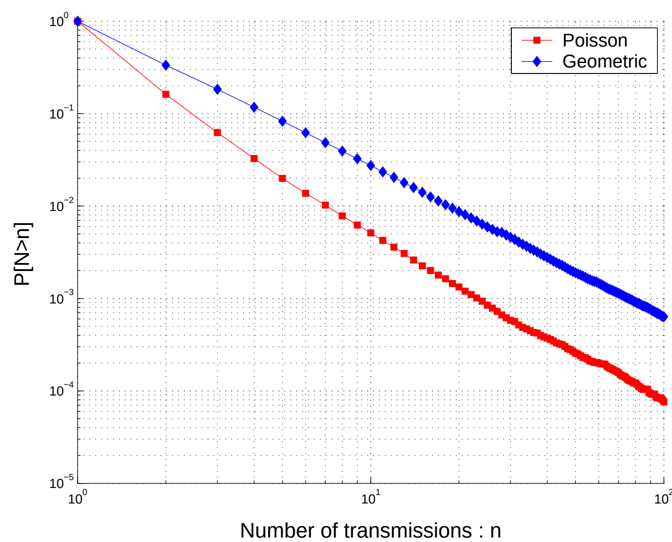

Fig. 6. Illustration for Example 4.

Example 4: Finally, as stated earlier in the paper, our results also hold for lattice, e.g., integer-valued random variables. The detailed proofs of these results will be written in the extended version of this paper [7]. Here, we justify this claim with two typical cases of discrete distributions, Poisson and geometric. For the case when A and L are Poisson, we select $\lambda_A = 8$ and $\lambda_L = 5$, respectively. The parameters of the geometric distributions for A and L are chosen as $p_A = 0.125$ and $p_L = 0.2$, respectively. The simulation results for 5×10^5 samples are illustrated in Figure 6 that clearly shows the linear characteristics of $\mathbb{P}[N > n]$ on the log-log plot, i.e., the power law behavior.

V. Engineering Implications

From an engineering perspective, our main discovery is the matching between the statistical characteristics of the channel and transmitted data (packets). Basically, if $\alpha \log \mathbb{P}[A > x] > \log \mathbb{P}[L > x]$ or $\alpha \log \mathbb{P}[A > x] < \log \mathbb{P}[L > x]$, then one can expect good or bad (measured by the existence of α-moment for N and T) delay performance. In this regard, we want to point out that our model depicted in Figure 1 admits a variant of a more general form which we term threshold crossing [7]. We believe that this model provides a basic structure that explains the power law phenomena in many natural and man-made systems. For example, we discover a new effect [11] that a basic finite population ALOHA model with variable size (exponential) packets is characterized by power law transmission delays, possibly even resulting in zero throughput. This power law effect might be diminished, or perhaps eliminated, by reducing the variability of packets. However, we show that even a slotted (synchronized) ALOHA with packets of constant size can exhibit power law delays when the number of active users is random.

Furthermore, on the physical layer, it is well known that wireless links, especially for low-powered sensor networks, have higher error rates than the wired counterparts. This may result in large delays on the data link layer due to the (IP) packet variability and channel failures. When the codewords, the basic units of packets in the physical layer, are much smaller than the maximum size of the packets, our results show that the number of retransmissions could be power law, which challenges the traditional model that assumes a geometric number of retransmissions. We believe that short codewords are realistic assumption for sensor networks, where complicated coding schemes are unlikely since the nodes have very limited computational power.

Since in reality, packet sizes may have an upper limit (e.g., WaveLAN's maximum transfer unit is 1500 bytes), this situation may result in truncated power law distributions for T and N in the main body with a stretched (exponentiated) support in relation to the support of L (see Example 3) and, thus, may result in very long, although, exponentially bounded delays. Similar investigations have been examined for truncated heavy-tailed distributions in the queueing context in [12].

Therefore, our results suggest that, packet fragmentation techniques need to be applied with special care since, if the packets are too small, they will mostly contain the packet header, which can limit the useful throughput; if packets are too large, power law delays can deteriorate the quality of transmission.

In conclusion, we would like to emphasize that, in practice, our results provide an easily computable benchmark for measuring the tradeoff between the data statistics and channel characteristics that permits/prevents satisfactory transmission.

VI. Proofs of Theorem 3 and Theorem 5

The proofs are based on large deviation results developed by Nagaev in [13]; specifically, we summarize Corollary 1.6 and Corollary 1.8 of [13] in this following lemma.

Lemma 2: Let $X_1, X_2, \cdots X_n$ and X be i.i.d random variables with $\mathbb{E}X = 0$, and $a_s^+ \triangleq \int_{u \ge 0} u^s d\mathbb{P}[X < u] < \infty$.

If $1 \le s \le 2$, then, for $x > y$ and $y^s \ge 4a_s^+$,

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i \ge x\right] \le n\mathbb{P}[X > y] + \left(\frac{ne^2 a_s^+}{xy^{s-1}}\right)^{x/2y}. \quad (26)$$

If $s > 2$, then,

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i \ge x\right] \le nc_s^{(1)} a_s^+ x^{-s} + \exp\left(\frac{-c_s^{(2)} x^2}{n\mathbb{V}ar[X_i]}\right), \quad (27)$$

where $c_s^{(1)} = (1 + 2/s)^s$, $c_s^{(2)} = 2(s+2)^{-2} e^{-s}$.

Proof: Please refer to [13]. ■

Now, we are ready to prove Theorem 3.

Proof: First, we establish the upper bound. For any $\delta > 0$, we have

$$\begin{aligned}\mathbb{P}[T > t] &= \mathbb{P}\left[\sum_{i=1}^{N-1} (U_i + A_i) + L > t\right] \\ &\le \mathbb{P}\left[\sum_{i=1}^{N} (U_i + A_i) > t, N \le \frac{t(1-\delta)}{\mathbb{E}[U + A]}\right] \\ &\quad + \mathbb{P}\left[N > \frac{t(1-\delta)}{\mathbb{E}[U + A]}\right] + \mathbb{P}[L > t] \\ &\triangleq I_1 + I_2 + I_3. \end{aligned} \quad (28)$$

In the following, we will show that $I_1 = o(1/t^\alpha)$ and $I_3 = o(1/t^\alpha)$. For I_3, the condition $\mathbb{E}[(U + A)^{1+\alpha+\epsilon}] < \infty$ and (2) imply $\mathbb{E}[L^{\alpha(1+\alpha+\theta)}] < \infty$, which renders

$$I_3 \le \frac{\mathbb{E}[L^{\alpha(1+\alpha+\theta)}]}{t^{\alpha(1+\alpha+\theta)}} = o\left(\frac{1}{t^\alpha}\right). \quad (29)$$

For I_1, let $X_i \triangleq (U_i + A_i) - \mathbb{E}[(U_i + A_i)]$ and $\zeta \triangleq (1 - \delta)/\mathbb{E}[U + A]$. We have $\mathbb{E}X_1^{1+\alpha+\delta} < \infty$, $\mathbb{E}X_1 = 0$ and

$$\begin{aligned} I_1 &\le \mathbb{P}\left[\sum_{i \le t(1-\delta)/\mathbb{E}[U+A]} (U_i + A_i) > t\right] \\ &= \mathbb{P}\left[\sum_{i \le \zeta t} X_i > \delta t\right]. \end{aligned} \quad (30)$$

Here we have two situations. If $1 < s = 1 + \alpha + \theta \le 2$, using (26) with $y = \delta t/2$, we obtain, as $t \to \infty$,

$$\begin{aligned}\mathbb{P}\left[\sum_{i\le\zeta t} X_i > \delta t\right] &\le \zeta t \mathbb{P}[X > \delta t/2] + \frac{e^2 2^{s-1} a_s^+}{t^{s-1}} \\ &\le \frac{\zeta 2^s \mathbb{E}[X^s]}{\delta^s t^{s-1}} + \frac{e^2 2^{s-1} a_s^+}{t^{s-1}} \\ &= o\left(\frac{1}{t^\alpha}\right). \end{aligned} \tag{31}$$

Otherwise, if $s = 1 + \alpha + \theta > 2$, by (27), we derive

$$\begin{aligned}\mathbb{P}\left[\sum_{i\le\zeta t} X_i > \delta t\right] &\le \zeta t \cdot c_s^{(1)} \frac{a_s^+}{(\delta t)^s} + \exp\left(\frac{-c_s^{(2)}(\delta t)^2}{\zeta t \mathbb{V}ar[X_i]}\right) \\ &= o\left(\frac{1}{t^\alpha}\right). \end{aligned} \tag{32}$$

Hence, from (30), (31) and (32), we derive

$$I_1 \le o\left(\frac{1}{t^\alpha}\right). \tag{33}$$

Recalling Theorem 2, we know

$$\lim_{t\to\infty} \frac{\log \mathbb{P}\left[N > \frac{t(1-\delta)}{\mathbb{E}[U+A]}\right]}{\log t} = -\alpha, \tag{34}$$

which, in conjunction with (29), (33) and (34), implies

$$\overline{\lim_{t\to\infty}} \frac{\log \mathbb{P}[T > t]}{\log t} \le -\alpha. \tag{35}$$

Next, we prove the lower bound. It is easy to obtain

$$\begin{aligned}\mathbb{P}[T > t] &= \mathbb{P}\left[\sum_{i=1}^{N-1}(U_i + A_i) + L > t\right] \\ &\ge \mathbb{P}\left[\sum_{i=1}^{N-1}(U_i + A_i) > t\right] \\ &\ge \mathbb{P}\left[\sum_{i=1}^{N-1}(U_i + A_i) > t, N \ge \frac{t(1+\delta)}{\mathbb{E}[U+A]} + 1\right] \\ &\ge \mathbb{P}\left[N \ge \frac{t(1+\delta)}{\mathbb{E}[U+A]} + 1\right] \\ &\quad - \mathbb{P}\left[\sum_{i=1}^{N-1}(U_i + A_i) \le t, N \ge \frac{t(1+\delta)}{\mathbb{E}[U+A]} + 1\right] \\ &\triangleq I_1 - I_2. \end{aligned} \tag{36}$$

Now, since

$$\begin{aligned} I_2 &\le \mathbb{P}\left[\sum_{i \le t(1+\delta)/\mathbb{E}[U+A]} (U_i + A_i) \le t\right] \\ &= \mathbb{P}\left[\sum_{i \le t(1+\delta)/\mathbb{E}[U+A]} (-X_i) \ge \delta t\right], \end{aligned}$$

and $(-X_i) \le \mathbb{E}[U + A] < \infty$, by using the same argument as in the proof for (33), we obtain

$$I_2 \le o\left(\frac{1}{t^\alpha}\right). \tag{37}$$

Again, by Theorem 2,

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

from which, using (36) and (37), we derive

$$\underline{\lim_{t\to\infty}} \frac{\log \mathbb{P}[T > t]}{\log t} \ge -\alpha. \tag{38}$$

Combining (35) and (38) completes the proof. ■

Next, we give a brief sketch of the proof of Theorem 5.

Proof: Using the same technique as in the proof of Theorem 3, we obtain

$$\mathbb{P}[T > t] \le \mathbb{P}\left[N \ge \frac{t(1-\delta)}{\mathbb{E}[U+A]}\right] + o\left(\frac{1}{t^\alpha}\right),$$

and

$$\mathbb{P}[T > t] \ge \mathbb{P}\left[N \ge \frac{t(1+\delta)}{\mathbb{E}[U+A]}\right] - o\left(\frac{1}{t^\alpha}\right).$$

By Theorem 4, we have

$$\mathbb{P}\left[N \ge \frac{t(1-\delta)}{\mathbb{E}[U+A]}\right] \sim \frac{\Gamma(\alpha+1)(\mathbb{E}[U+A])^\alpha}{(t(1-\delta))^\alpha \cdot H(\alpha \log t)^\alpha}, \tag{39}$$

and

$$\mathbb{P}\left[N \ge \frac{t(1+\delta)}{\mathbb{E}[U+A]}\right] \sim \frac{\Gamma(\alpha+1)(\mathbb{E}[U+A])^\alpha}{(t(1+\delta))^\alpha \cdot H(\alpha \log t)^\alpha}. \tag{40}$$

Finally, passing δ to zero in (39) and (40) completes the proof. ■

REFERENCES

[1] P. M. Fiorini, R. Sheahan, and L. Lipsky, "On unreliable computing systems when heavy-tails appear as a result of the recovery procedure," *ACM SIGMETRICS Performance Evaluation Review*, vol. 33, no. 2, pp. 15–17, September 2005.

[2] R. Sheahan, L. Lipsky, P. Fiorini, and S. Asmussen, "On the completion time distribution for tasks that must restart from the beginning if a failure occurs," *MAMA 2006 Workshop, Saint-Malo, France*, June 2006.

[3] D. P. Bertsekas and R. Gallager, *Data Networks*, 2nd ed. Prentice Hall, 1992.

[4] T. S. Rappaport, *Wireless communications: principles and practise*, 2nd ed. Prentice Hall, January 2002.

[5] V. Kulkarni, V. Nicola, and K. Trivedi, "The completion time of a job on a multimode system," *Advances in Applied Probability*, vol. 19, no. 4, pp. 932–954, Dec. 1987.

[6] K. Park and W. Willinger, Eds., *Self-Similar Network Traffic and Performance Evaluation*. A Wiley-Interscience Publication, Sept. 2000.

[7] P. R. Jelenković and J. Tan, "Threshold crossing and heavy tails," Dept. of Electrical Engineering, Columbia University, Tech. Rep., 2007, in preparation.

[8] S. M. Ross, *A First Course in Probability*, 6th ed. Prentice Hall, 2002.

[9] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation*, G.-C. Rota, Ed. Cambridge University Press, 1987, vol. 27.

[10] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis: Queues, Communications, and Computing*. Chapman and Hall, London, UK, 1995.

[11] P. R. Jelenković and J. Tan, "Is ALOHA causing power law delays?" Dept. of Electrical Engineering, Columbia University, Tech. Rep., August 2006, in submission.

[12] P. R. Jelenković, "Network multiplexer with truncated heavy-tailed arrival streams," *INFOCOM'99, New York, NY*, March 1999.

[13] S. V. Nagaev, "Large deviations of sums of independent random variables," *The Annals of Probability*, vol. 7, no. 5, pp. 745–789, 1979.