Retransmission Delays with Bounded Packets: Power law body and Exponential tail

Jian Tan, B. T. Swapna, and Ness B. Shroff
Departments of ECE and CSE
The Ohio State University

***Abstract*—Retransmissions serve as the basic building block that communication protocols use to achieve reliable data transfer. Until recently, the number of retransmissions were thought to follow a geometric (light-tailed) distribution. However, recent work shows that when the distribution of the packet sizes have infinite support, retransmission-based protocols may result in heavy tailed delays and possibly zero throughput even when the afore-mentioned distribution is light-tailed. In reality, however, packet sizes are often bounded by the Maximum Transmission Unit (MTU), and thus the aforementioned result merits a deeper investigation.**

To that end, in this paper, we allow the distribution of the packet size L to have finite support. Under mild conditions, we show that the transmission duration distribution exhibits a transition from a power law main body to an exponential tail. The time scale to observe the power law main body is roughly equal to the average transmission duration of the longest packet. The power law main body, if significant, may cause the channel throughput to be very close to zero. These theoretical findings provide an understanding on why some empirical measurements suggest heavy tails. We use these results to further highlight the engineering implications of distributions with power law main bodies and light tails by analyzing two cases: (1) The throughput of on-off channels with retransmissions, where we show that even when packet sizes have small means and bounded support the variability in their sizes can greatly impact system performance. (2) The distribution of the number of jobs in an $M/M/\infty$ queue with server failures. Here we show that retransmissions can cause long-range dependence and quantify the impact of the maximum job sizes on the long-range dependence.

***Index Terms*—Retransmissions, delay distribution, throughput, power law body, heavy tails.**

I. INTRODUCTION

Retransmissions are fundamental in ensuring reliable data transfer over communication networks with channel errors. Traditionally, retransmissions were assumed to result in light-tailed (rapidly decaying tail distribution) transmission delays. The conventional belief was that the number of retransmissions follows a geometric distribution [1], which is true when the errors are independent of the size of the transmitting packet. However, recent work [2]–[4] shows that when the probability of packet errors is an increasing function of the packet length, which is often true in communication networks, the number of retransmissions do not follow a geometric distribution. The rough intuition is as follows: If we use the traditional retransmission schemes that repeatedly send a packet until it is received successfully, the expected transmission duration of sending an $N-$bit packet over an *i.i.d.* binary erasure channel grows on the order $O(1/p^N)$, where $1-p$ is the per-bit erasure probability. Since the expected transmission duration grows exponentially in the number of bits N, even with light-tailed packet sizes, where the distribution of N decreases at least exponentially fast, the resultant delay is still heavy-tailed. In fact, it has been shown in [2]–[4], under the assumption that the packet size distribution has infinite support, that all retransmission-based protocols could cause heavy-tailed behavior (specifically, power law transmission durations) and possibly even zero throughput, even when the data units and channel characteristics are light-tailed. Following this observation, there have been several efforts to identify transmission schemes to mitigate the power law delays. In [5], the authors show that independent or bounded fragmentation guarantees light-tailed completion time as long as the packet/file size is light-tailed. This scheme requires additional overhead for each packet transmission, hence resulting in significant throughput loss. In [6], the authors consider the use of fixed-rate coding techniques to transmit information in order to mitigate delays. Their study reveals a complicated relationship between the coding complexity and the transmission delay/throughput. They characterize the possibility of transmission delay following a power law with index less than one when the coding complexity is high and when the receiver does not have a memory of successfully received bits. In [7], the authors investigate the use of multi-path transmission schemes such as redundant, and split transmission techniques. They find that the power-law transmission delay phenomenon still persists with multi-path transmission under the assumption that the packet size distribution has infinite support. However, in practice, packet sizes are bounded by the maximum transmission unit (MTU). This fact motivates us to more carefully investigate the impact that retransmissions have on network performance by allowing the packet sizes to have finite support.

We consider a system where the channel dynamics are modeled by an *on-off* process $\{(A_i, U_i)\}_{i\geq 1}$ where A_i corresponds to the time when the channel is available and U_i the time period when the channel is not available, as in [8]. Let L be the random variable that denotes the length of a generic packet. At the beginning of each available period A_i, we attempt to transmit the packet. If $L < A_i$, we say that the transmission is successful; otherwise, we wait until the beginning of the next available period A_{i+1} and retransmit the packet from the

This material is based upon work supported in part by, the U.S. Army Research Laboratory and the U. S. Army Research Office under grant number W911NF-08-1-0238 and by the National Science Foundation through grants CNS-1065136 and CNS-1012700.

beginning. As mentioned earlier, we focus on the situation of practical interest, i.e., when the distribution of L has finite support on the interval $[0, b]$. We study the asymptotic properties of the distributions of the total transmission time and number of retransmissions. Our main contributions in this paper can be summarized as follows:

(I) Under a general polynomial relationship between the packet size distribution and channel available period distribution (this relationship provides a measure of the quality of the channel), we show that, even when the packet size has an upper limit, the transmission duration distribution is characterized by a power law main body. This power law behavior spans over a time scale that is approximately equal to the average transmission duration of the longest packet. Additionally, we show that this distribution eventually becomes light-tailed. We characterize the transition of the transmission delay distribution from a power law main body to an exponential tail. Thus, depending on the probabilities of interest and the system parameters, the transmission delays may experience heavy or light-tailed distributions. More importantly, both the power law main body and the exponential tail could dominate performance. When this power law main body is significant, it could possibly cause the channel throughput to be very close to zero (as shown in Theorem 4), implying that some careful re-examination and adjustment of system parameters are needed. On the other hand, if the exponential tail is more evident, this suggests that the system is operating in a benign environment. Similar phenomenon of power law up to a certain threshold followed by an exponential decay has been observed for inter-contact time distributions between mobile devices [9].

(II) Using the afore-mentioned results, we study two cases of interest. First, we investigate the system throughput when the packet lengths have an upper limit b. Our results show that under certain conditions the channel throughput may be very close to zero for large b even when the average packet size is very small. Next, we study an $M/M/\infty$ queue with server failures. When active servers fail according to *i.i.d.* Poisson point processes, we observe that the number of jobs in the system exhibits long-range dependence. This effect can be eliminated if job sizes are upper bounded. However, we find that there may still be a strong autocorrelation for the number of jobs in the system that spans over a large time interval for bounded job sizes, implying that the system may exhibit long-range dependence over operating regions of interest.

These theoretical findings provide a new understanding on the controversy in empirical measurements why heavy tails are observed for certain measurements and light tails for others (e.g., wireless networks). The discovery of heavy-tailed statistical characteristics of traffic streams in modern computer networks [10] led to an extensive amount of research on the issue of power laws in information networks.

For example, it was suggested in [11] that the transmission delay distribution in IEEE 802.11 wireless ad hoc networks can be expressed as a power law. There are also different views advocating other distributions as providing the correct description of the system. For example, the empirical measurement in [12] shows that link delays over the wireless mesh network are fitted by either gamma or logistic distributions. These seemingly contradicting results have been addressed in [13], which suggests that (1) some claims on the heavy/light tails may not be legitimate due to the lack of sufficient measurements for the hypothesis testing, and (2) engineers should focus on the behavior of a distribution's "waist" that refers to the portion for which there are enough data to summarize the distributional information. Our results provide the mathematical basis for understanding these competing claims and show that indeed depending on the operating points and parameters of interest, either heavy or light tail phenomenon may dominate performance.

Also, from an engineering perspective, our results further emphasize the insight developed in [8] that retransmissions may significantly amplify the packet size variability to much larger variability in transmission delays. More precisely, if there is a polynomial functional relationship between the distributions of the channel ON periods and the packet size, the transmission duration is very close to a power law distribution over the time scale of order $1/\mathbb{P}[A > b]$. Thus, even for packet lengths with small MTUs, the small variability in the packet size can still be amplified by the retransmission based protocols, causing potentially poor performance. This observation could be a possible explanation for the empirical measurements in [14], which claims that the utilization of the 802.11 protocol is lower due to retransmissions.

Our results also provide insights in designing control algorithms. For example, in physical layer, power control can change the rate at which the packet is transmitted. In this sense, power control can be thought as a way to change the relationship between the channel dynamics and the units in which packets should be transmitted in order to achieve the best network performance.

II. Model description and Related work

In this section, we formally describe our model and provide necessary definitions and notation. Some related results are also presented in this part.

Throughout this paper, a positive function f is called regularly varying (at infinity) with index ρ if $\lim_{x\to\infty} f(\lambda x)/f(x) = \lambda^\rho$ for all $\lambda > 0$. It is called slowly varying if $\rho = 0$ [15]. For any two real functions $f(t)$ and $g(t)$ we use $f(t) \sim g(t)$ as $t \to \infty$ to denote $\lim_{t\to\infty} f(t)/g(t) = 1$. Similarly, we say that $f(t) \gtrsim g(t)$ as $t \to \infty$ if $\liminf_{t\to\infty} f(t)/g(t) \geq 1$; $f(t) \lesssim g(t)$ has a complementary definition. We use "$\overset{d}{=}$" and "$\overset{d}{\leq}(\overset{d}{\geq})$" to denote equal in distribution and less (greater) than or equal in distribution, respectively. We use $\vee$ to denote $\max$, i.e., $x \vee y \equiv \max\{x, y\}$.

In this paper, we adopt the retransmission model that was proposed in [8]. The channel dynamics are modeled as an on-off process $\{(A_i, U_i)\}_{i\geq 1}$ that alternates between available

A_i and unavailable U_i periods, respectively. Let L denote the random length of a generic packet. At the beginning of each time period A_i when the channel becomes available, we attempt to transmit the packet. If $L < A_i$, we say that the transmission is successful; otherwise, we wait until the beginning of the next available period A_{i+1} and retransmit the packet from the beginning. This process continues until the packet is successfully transmitted over the channel. In this paper, we assume that $\{U_i\}_{i\geq 1}$ and $\{A_i\}_{i\geq 1}$ are two mutually independent sequences of *i.i.d.* random variables with $U_i \stackrel{d}{=} U$, $A_i \stackrel{d}{=} A$ and U independent of A. A sketch of the model depicting the system is drawn in Figure 1; see also in [8]. As

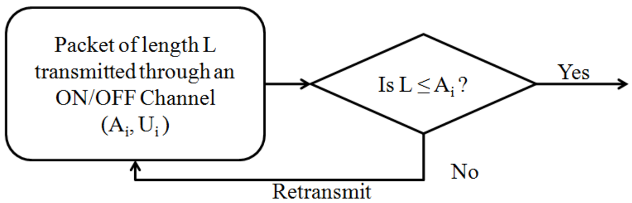

Fig. 1. Packets sent over channels with failures

mentioned earlier, unlike [8], we allow the packet length L to take values on finite interval $[0, b], b > 0$. Our goal will be to study the behavior of the number of retransmissions $N(b)$ and the total transmission delay $T(b)$ as b scales with the number of retransmissions.

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

$$N(b) \triangleq \inf\{n : A_n > L\},$$

and, the total transmission time for the packet is defined as

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

First, we state a condition on the moment generating function of U that will be used in the characterization the total transmission time $T(b)$.

Condition 1: $\mathbb{E}\left[e^{\theta U}\right] < \infty$ for all $\theta > 0$.

This condition implies that U has a light tailed distribution. Note that this assumption makes the problem non-trivial. Because if U were heavy tailed, it would not at all be surprising that the transmission delay would be heavy tailed.

Next, we briefly discuss the results from [8] that investigate the heavy-tailed (specifically, power law) delay behavior of retransmission based protocols. When the distribution function of L has an infinite support ($b = \infty$), it has been shown in Lemma 1 of [8] and Proposition 1.2 of [4] that both the transmission time and the number of transmissions follow subexponential distributions, as given by the following result.

Proposition 1 (from [8]): If $\mathbb{P}[L > x] > 0$ for all $x \geq 0$, then both $N(\infty)$ and $T(\infty)$ are subexponential in the following sense that, for any $\epsilon > 0$, as $n \to \infty$ and $t \to \infty$,

$$e^{\epsilon n}\mathbb{P}[N(\infty) > n] \to \infty, \quad e^{\epsilon t}\mathbb{P}[T(\infty) > t] \to \infty.$$

This class of heavy-tailed distributions has a rich structure, including power laws, heavy-tailed Weibull distributions and nearly exponential distributions. For a detailed analysis of this class of distributions induced by retransmissions see [2]. Since power law distributions are closely related to long range dependency (see Section V-B) and channel stability (see Section V-A), we focus on power law delays. Below we quote Theorem 2 of [8] as Proposition 2 (see also Theorem 2.5 in [2] and Theorem 2.2 in [4]), which show that both $N(\infty)$ and $T(\infty)$ can follow power law distributions.

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

$$\lim_{x\to\infty} \frac{\log \mathbb{P}[L > x]}{\log \mathbb{P}[A > x]} = \alpha,$$

then,

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

Additionally, if $\mathbb{E}\left[e^{\theta(A+U)}\right] < \infty$ for some $\theta > 0$, then,

$$\lim_{t\to\infty} \frac{\log \mathbb{P}[T(\infty) > t]}{\log t} = -\alpha. \tag{2}$$

Remark 1: In order to prove Equation (2), we only need a weaker condition $\mathbb{E}\left[U^{(\alpha\vee 1)+\theta}\right] < \infty$ and $\mathbb{E}\left[A^{1+\theta}\right] < \infty$ for some $\theta > 0$. See [2] for more details.

III. Transition from heavy to light tails

In this section, we present preliminary investigations and a motivating example when the packet sizes have finite support. When b is finite, unlike the case of infinite support that can cause subexponential delays, both the number of retransmissions $N(b)$ and the total transmission time $T(b)$ have exponential tails, as shown in the following proposition. When $U_i \equiv 0$ and A has a density function, this case has been studied in Corollary 3.1 of [4].

Proposition 3: Under Condition 1, for $b < \infty$ with $0 < \mathbb{P}[A \leq b] < 1$, if $\mathbb{P}[A + U > x | A \leq b]e^{\gamma^* x}$ is nonlattice and directly Riemann integrable with $\gamma = \gamma^*$ being the solution of $\int_0^\infty e^{\gamma s} d\mathbb{P}[A + U \leq s | A \leq b] = 1/\mathbb{P}[A \leq b]$, then

$$\mathbb{P}[N(b) > n] \leq (\mathbb{P}[A \leq b])^n, \quad \mathbb{P}[T(b) > t] \lesssim Ce^{-\gamma^* t},$$

where

$$C = \frac{e^{\gamma^* b}\mathbb{P}[A > b]\mathbb{P}[A \leq b]^{-1}}{\gamma^* \int_0^\infty s e^{\gamma^* s} d\mathbb{P}\left[A + U \leq s | A \leq b\right]}.$$

Proof: See Appendix. ■

Remark 2: Note that if the packet size has bounded support, the number of retransmissions being geometrically distributed immediately follows. The other conditions in Proposition 3 are needed to characterize the tail of the transmission delay distribution.

Heavy tails and light tails have very different statistical characteristics. The preceding cases represent the two extremes that feature heavy tails and light tails, corresponding to a fixed finite and infinite b, respectively. It motivates us to study the transition from heavy tails to light tails when b scales such that it is neither a fixed finite constant nor equal to infinity. To that end, we introduce a hidden random variable L^* that has an infinite support, and the packet size L satisfies the following condition

$$\mathbb{P}[L > x] = \begin{cases} \mathbb{P}[L^* > x | L^* \le b] & x \le b \\ 0 & x > b, \end{cases} \tag{3}$$

where $\mathbb{P}[L^* \le b] > 0$. Clearly, when b changes, the distribution of L changes accordingly with respect to b. Thus, we also use the notation $L_b \equiv L$ for increased clarity when necessary.

We use the following numerical evaluation to further illustrate the effect of transition from the power law main body to the light tail for the transmission delay distribution.

Example 1: This numerical example shows that the delay distribution has a power law main body and an exponential tail when L takes values on a finite interval $[0, b]$. Assume that $U_i = 0$ and both L^* and A follow exponential distributions with rate 0.8 and 1.0, respectively. We plot $\mathbb{P}[T(b) > t]$ on the log-log scale by changing b from 4 to 10. It is clear from the figure that the support of the power law main body increases very quickly with respect to b: a small increment of b from 4 to 10 results in a big expansion of the power law support from nearly 10 to 1000, which is a dramatic amplification.

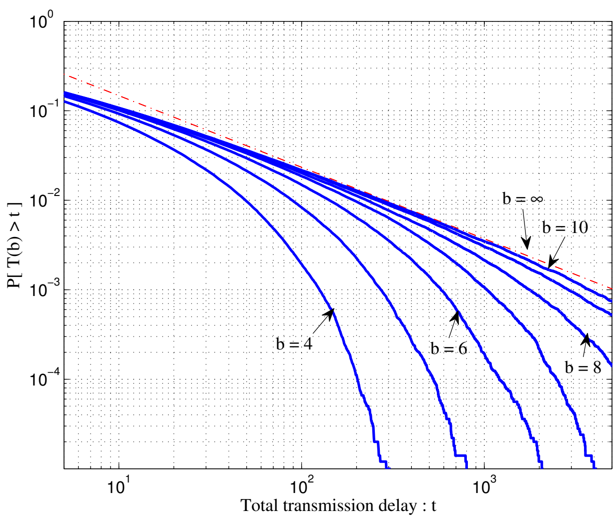

Fig. 2. The transmission delay distribution. A small increase in the maximum packet length, b, results in a large expansion of the power law main body.

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

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

Lemma 1: If $0 < b_1 \le b_2$, then $N(b_1) \overset{d}{\le} N(b_2)$ and $T(b_1) \overset{d}{\le} T(b_2)$.

Remark 3: This result is supported by the preceding example. As easily seen from Figure 2, $\mathbb{P}[T(b) > t]$ is monotonically increasing as b increases.

Proof: First, we want to show that $L_{b_1} \overset{d}{\le} L_{b_2}$. Recalling (3), it is sufficient to show that, for $x \ge 0$,

$$\frac{\mathbb{P}[x < L^* \le b_1]}{\mathbb{P}[L^* \le b_1]} \le \frac{\mathbb{P}[x < L^* \le b_2]}{\mathbb{P}[L^* \le b_2]}.$$

It is easy to verify the preceding inequality by checking

$$(\mathbb{P}[L^* \le b_1] - \mathbb{P}[L^* \le x])\,\mathbb{P}[L^* \le b_2] \le$$
$$(\mathbb{P}[L^* \le b_2] - \mathbb{P}[L^* \le x])\,\mathbb{P}[L^* \le b_1].$$

Now, by Definition 1, both $N(b)$ and $T(b)$ is monotonically increasing in L. Therefore, we prove that $N(b_1) \overset{d}{\le} N(b_2)$ and $T(b_1) \overset{d}{\le} T(b_2)$. ∎

IV. Main Results

This section presents our main results. Here, we assume that $\bar{F}(x)$ is a continuous function with support on $[0, \infty)$; note that L has a finite support on $[0, b]$ by Equation (3). We first state a condition that characterizes the functional relationship between the distributions of packet sizes and channel available periods.

Condition 2: $\lim_{x \to \infty} \frac{\log \mathbb{P}[L^* > x]}{\log \mathbb{P}[A > x]} = \alpha.$

We now present our main results that characterize the transition of the distributions of $N(b)$ and $T(b)$ from power law main bodies to exponential tails. Note that $\left(\bar{G}(b)\right)^{-1} = \left(\mathbb{P}[A > b]\right)^{-1}$ is the expected number of transmissions for a packet of size b.

Theorem 1: If Condition 2 holds for some $\alpha > 0$, then, for fixed $0 < \eta < 1$, and $\epsilon > 0$, there exist n_0 and b_0 such that for any $b > b_0$, we have, for $n_0 < n < \left(\bar{G}(b)\right)^{-\eta}$,

$$1 - \epsilon < \frac{\log \mathbb{P}\left[N\left(b\right) > n\right]}{-\alpha \log n} < 1 + \epsilon, \tag{4}$$

In addition, if $\bar{G}(x)$ is left-continuous at

$$b^* \triangleq \sup\left\{x : \mathbb{P}[L > x] > 0\right\} > 0,$$

we have,

$$\lim_{n \to \infty} \frac{\log \mathbb{P}\left[N\left(b\right) > n\right]}{n} = \log\left(1 - \bar{G}(b^*)\right). \tag{5}$$

Remark 4: This result shows that when b is large, albeit finite, the distributions of $N(b)$ and $T(b)$ consist of two different regions: a power law main body, and an exponential tail. Equation (4) implies that $\mathbb{P}[N(b) > n]$ is approximately a power law distribution with index α when n is smaller than $\left(\bar{G}(b)\right)^{-1}$, which is the average number of retransmissions of the largest packet. Equation (5) suggests an exponential tail distribution when n is large.

Proof: See Appendix. ∎

While the detailed proof is given in the Appendix, we provide a brief outline here. The upper bound in (4) follows from Lemma 1, and Proposition 2. Regarding the lower bound, we use the relation $\mathbb{P}(N(b) > n) = \mathbb{E}[(1 - \bar{G}(L))^n]$ along with the Condition 2, and $n < (\bar{G}(b))^{-\eta}$ to obtain a tight lower bound on $\mathbb{P}[N(b) > n]$ of the order $1/n^{\alpha - \epsilon}$ for any small $\epsilon > 0$. This characterizes the region of interest that is well approximated by power law distributions for the number of retransmissions. The bounds in (5) are also obtained by using the relation $\mathbb{P}(N(b) > n) = \mathbb{E}[(1 - \bar{G}(L))^n]$. The upper bound uses the fact that $\mathbb{P}[L \le b^*] = 1$. The lower bound is obtained by truncating the expectation in the interval $[b^* - \epsilon, b^*]$ and by using the fact that the $\mathbb{P}[b^* - \epsilon < L < b^*] > 0$. Continuity condition on $\bar{G}(x)$ allows us to obtain the result in (5) by letting $\epsilon \to 0$.

Theorem 2: If Condition 2 holds for some $\alpha > 0$ along with $\mathbb{E}\left[e^{\theta(A+U)}\right] < \infty$ for some $\theta > 0$, then, for fixed $0 < \eta < 1$, and $\epsilon > 0$, there exist t_0 and b_0 such that for any $b > b_0$, we have, for $t_0 < t < \left(\bar{G}(b)\right)^{-\eta}$,

$$1 - \epsilon < \frac{\log \mathbb{P}\left[T\left(b\right) > t\right]}{-\alpha \log t} < 1 + \epsilon. \tag{6}$$

In addition, if $\bar{G}(x)$ is left-continuous at $b^* = \sup\{x : \mathbb{P}[L > x] > 0\} > 0$, we have,

$$\lim_{t\to\infty} \frac{\log \mathbb{P}\left[T\left(b\right) > t\right]}{t} = -\gamma^*, \tag{7}$$

where $\gamma = \gamma^*$ is the solution of

$$\int_0^\infty e^{\gamma s} d\mathbb{P}\left[A + U \le s | A \le b^*\right] = 1/(1 - \bar{G}(b^*)). \tag{8}$$

Proof: See Appendix. ■

Under a more restrictive condition, we can obtain the following more precise result that characterizes the exact asymptote for the power law main body.

Theorem 3: If

$$\mathbb{P}[L^* > x]^{-1} \sim \Phi\left(\mathbb{P}[A > x]^{-1}\right), \tag{9}$$

where $\Phi(\cdot)$ is regularly varying with index $\alpha > 0$, then, for $\epsilon > 0$, there exist $\zeta, n_0, t_0 > 0$ and b_0 such that for any $b > b_0$, we have

$$\sup_{n_0 < n < \zeta/\bar{G}(b)} \left| \frac{\mathbb{P}\left[N\left(b\right) > n\right]\Phi(n)}{\Gamma(\alpha+1)} - 1 \right| < \epsilon, \tag{10}$$

where $\Gamma(\cdot)$ is the gamma function. Additionally, if $\mathbb{E}\left[e^{\theta(A+U)}\right] < \infty$ for some $\theta > 0$, then,

$$\sup_{t_0 < t < \zeta/\bar{G}(b)} \left| \frac{\mathbb{P}\left[T\left(b\right) > t\right]\Phi(t)}{\Gamma(\alpha+1)} - 1 \right| < \epsilon. \tag{11}$$

Proof: See Appendix. ■

Remark 5: The preceding results imply Equations (4) and (6), and characterize the exact asymptote for the power law main body.

The proof of Theorem 3 combines the proof of Theorem 2.7 in [2] and the same techniques used in proving Theorems 1 and 2 above. Given the characterization of distribution of L^* in terms of A as in (9), we make use of the Characterization Theorem of regular variation and the uniform convergence of slowly varying functions to obtain uniform bounds in (10). Next, we state a lemma that bounds the distributions of $N(b), T(b)$ when the number of retransmissions and transmission time are between the power law main body and the exponential tail.

Lemma 2: If Condition (2) holds, then for any fixed $\eta_2 > \eta_1 > 1$ and $0 < \epsilon < \eta_1$, there exists b_0, such that for all $b > b_0$, and for $\left(\bar{G}(b)\right)^{-\eta_1} < n < \left(\bar{G}(b)\right)^{-\eta_2}$,

$$-n^{1-\frac{1}{\eta_2+\epsilon}} \le \log \mathbb{P}\left[N\left(b\right) > n\right] \le -n^{1-\frac{1}{\eta_1-\epsilon}} \tag{12}$$

In addition, if $\mathbb{E}\left[e^{\theta(A+U)}\right] < \infty$ for some $\theta > 0$, then for $\left(\bar{G}(b)\right)^{-\eta_1} < t < \left(\bar{G}(b)\right)^{-\eta_2}$, we have

$$-t^{1-\frac{1}{\eta_2+\epsilon}} \le \log \mathbb{P}\left[T\left(b\right) > t\right] \le -t^{1-\frac{1}{\eta_1-\epsilon}}. \tag{13}$$

Proof: See Appendix. ■

Loosely speaking, Lemma 2 characterizes the transition between the power law main body and the exponential tail.

A. Numerical Evaluation

We next present numerical results to support our results in this section. First, we verify the support region of the power law main body that is characterized by Theorem 2.

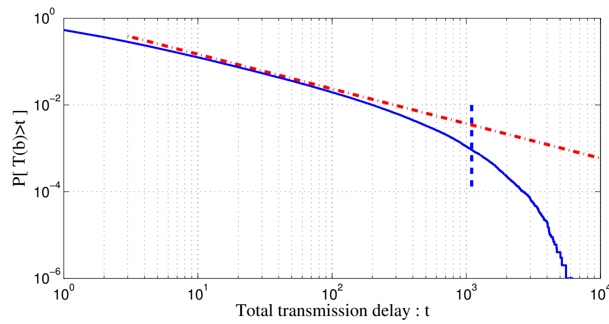

Fig. 3. Power law main body is close to the asymptote computed using Theorem 3 on the interval corresponding to the average transmission time for a packet of size b.

Under the setting of Example 1, we plot the distribution of the total transmission delay $T(b)$ for $b = 7$ in Figure 3. The dotted line represents the asymptote $\Gamma(1.0/0.8)t^{-1.0/0.8}$ that is computed using (11). The dashed vertical line corresponds to the average transmission time for a packet of size $b = 7$ over the channel, i.e., $\mathbb{E}[A]/\bar{G}(b) = e^7 = 1.096 \times 10^3$. It is easy to see a power law main body that is close to the computed asymptote on the interval $[0, 1.096 \times 10^3]$.

Next, we investigate the exponential tail under the setting of Example 1. We plot the distribution of the number of retransmissions $N(b)$ for $b = 3$ in Figure 4. It can be observed that, after taking logarithm with base 10, the tail distribution is asymptotically a straight line. The slope of the line is -0.022, which matches our theoretical result $\log_{10}\left(1 - \bar{G}(b)\right) = \log_{10}\left(1 - e^{-3}\right)$ that is computed by equation (5).

Fig. 4. Exponential tail for the distribution of the number of retransmissions.

V. Related models

In this section, using the results obtained in Section IV, we study two related models that further highlight the engineering implications. In Section V-A, we show that, even when the packet has a small mean with a bound that does not depend on its upper limit b, due to the power law main body caused by the polynomial functional relationship between the packet size variability and channel dynamics, the channel throughput may still be very close to zero. Therefore, both choosing an appropriate upper limit b and designing the packet variability are important for the system performance. In Section V-B, we investigate the autocorrelation function of the number of jobs in a $M/M/\infty$ queue with server failures. If a server failure occurs during the processing of a job, this job has to restart from the beginning on the same server. We show that, under the condition that the job sizes follow exponential distribution (with infinite support), this model can cause long range dependence. Introducing an upper limit to all job sizes can eliminate the long range dependence; however, the strong dependence can span over a large interval, which implies a performance very close to long range dependence.

A. Throughput of the on-off channel

Consider the same on-off channel model as in Section II. Now, suppose that the source has an infinite number of backlogged packets to be sent. Let $\{L_i\}_{i\geq 1}$ be the *i.i.d.* sequence of packet sizes. Define $\{T_i\}$ to be the duration for transmitting packet L_i. Note that the channel is still available immediately after the successful transmission of packet L_i, thus we can start transmitting L_{i+1} without waiting for the next available period. Due to this effect, the durations $\{T_i\}$ are not independent random variables. We are interested in studying the throughput $\Lambda_n(b)$ of this system for the first n packets,

$$\Lambda_n(b) \triangleq \frac{\sum_{i=1}^n L_i}{\sum_{i=1}^n T_i}.$$

Since it is not clear whether $\Lambda_n(b)$ converges to a limit as n goes to infinity, we use $\overline{\underline{\lim}}$ to denote both $\limsup$ and $\liminf$.

Theorem 4: If Condition 2 holds for some $0 < \alpha < 1$ and $\mathbb{E}[e^{\theta(A+U)}] < \infty$ for some $\theta > 0$, then, the throughput of the on-off channel $\Lambda_n(b)$ satisfies, as $b \to \infty$,

$$\overline{\underline{\lim}}_{n\to\infty} \frac{\log \Lambda_n(b)}{\log \bar{G}(b)} \sim (1-\alpha). \tag{14}$$

Proof: See Appendix. ■

Remark 6: From the preceding result, if both L^* and A follows exponential distributions with rate μ and ν, respectively, then, as the upper limit b goes to ∞, the throughput vanishes to zero roughly with speed $e^{-(\nu-\mu)b}$ for $\nu > \mu$. Figure 5 shows this exponential rate of decay of throughput for the on-off channel with $\mu = 0.5$ and $\nu = 1$. Note that the mean packet size $\mathbb{E}[L]$ is bounded by a constant $\mathbb{E}[L^*] = 1/\mu$ that does not depend on b. This suggests that some special care needs to be taken when engineering retransmission based protocols: both the maximum packet size and the variability of the packet size can impact the throughput.

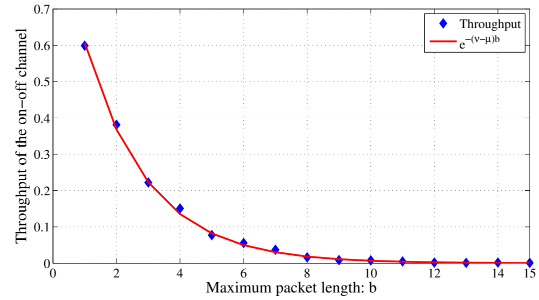

Fig. 5. Throughput of the on-off channel for a finite b. L^* and A follow exponential distributions of rate $\mu = 0.5$ and $\nu = 1$.

B. Long-range dependence in M/M/∞ queues with failures

Consider a $M/M/\infty$ queue with Poisson arrival rate λ. Let L^* be an exponential random variable with rate μ. Suppose that all the active servers fail independently according to Poisson point processes with the same rate ν. Immediately after a server fails in the middle of processing a job that runs on this server, this job has to restart on the same server from the beginning. Let T_i denote the processing time for job i. Let $M(t)$ be the number of jobs in the system at time t. Theorem 5 shows that $M(t)$ may be long range dependent if all the job sizes $\{L_i\}_{i>-\infty}$ follow an exponential distribution with rate μ; $L_i \overset{d}{=} L^*$. In the rest of this section, we assume that the system has reached stationarity.

Theorem 5: If $1 < \mu/\nu < 2$, then $M(t)$ is long range dependent in the sense that

$$\int_0^\infty \mathrm{Cov}(M(t), M(t+s))ds = \infty.$$

Remark 7: If $\mu/\nu \leq 1$, then $\mathbb{V}ar[M(t)] = \infty$, and $\mathrm{Cov}(M(t), M(t+s))$ is not defined for this case. If $\mu/\nu = 2$, the result in this theorem also holds, but we omit the proof.

Proof: By the well-known result on the autocorrelation function of the number of jobs in a $M/G/\infty$ queue (e.g., see [16]), we obtain

$$\mathrm{Cov}(M(t), M(t+s)) = \lambda \int_s^\infty \mathbb{P}[T_1 > x]dx.$$

Proposition 2 implies that, for $0 < \epsilon < 2 - \mu/\nu$,

$$\mathbb{P}[T_1 > x] \gtrsim \frac{1}{x^{\mu/\nu+\epsilon}},$$

as $x \to \infty$. Therefore, we obtain, as $s \to \infty$,

$$\mathrm{Cov}(M(t), M(t+s)) \gtrsim \frac{\lambda}{s^{\mu/\nu-1+\epsilon}},$$

which, using the fact that $0 < \mu/\nu - 1 + \epsilon < 1$, completes our proof. ■

This long range dependence can be eliminated by restricting the support of the job size distribution. In the following, similar to Section IV, we assume that all the *i.i.d.* job sizes $\{L_i\}$ follow a truncated exponential distribution with $\mathbb{P}[L_1 \geq x] = \mathbb{P}[L^* \geq x | L^* \leq b]$ for $b > 0$. We show in Theorem 6 that the integration of the autocorrelation function may still be very large, indicating a performance close to long-range dependence.

Theorem 6: If $1 < \mu/\nu < 2$ and job sizes have a finite support on $[0, b]$, then, as $b \to \infty$,

$$\log \int_0^{\infty} \mathrm{Cov}(M(t), M(t+s))ds \sim (2\nu - \mu)b,$$

which goes to infinity linearly as $b \to \infty$.

Proof: See Appendix. ■

To prove the asymptotic relationship in Theorem 6, we show that $(2\nu - \mu)b$ is both an asymptotic upper bound as well as a lower bound to the logarithm of the covariance. To show the lower bound, we make use of the lower bound in (6). To prove the upper bound, we appeal to the upper bound in Proposition 2 and Lemma 2.

VI. Conclusion

In practical communication protocols, all packets are bounded by an upper limit, say, the maximum transmission unit (MTU). Our results show that, for retransmission mechanism, this upper limit has an important influence on the system performance.

We show that the retransmission mechanism could enlarge transmission durations in a highly non-linear manner. Under a general polynomial relationship between the statistical characteristics of channel dynamics and packet size variability, i.e., $\log \mathbb{P}[L > x] \approx \alpha \log \mathbb{P}[A_1 > x], x \le b$, where b is the maximum packet length, the time for a successful transmission approximately follows a power law over the time scale of order $1/\mathbb{P}[A > b]$, i.e., the average transmission time of the longest packet. Thus, even for packets with an upper limit, a small variability in the packet size distribution can still be amplified dramatically by retransmission based protocols, possibly causing very poor performance if α is small, e.g., $\alpha < 1$.

These effects could greatly impact the system performance in many engineering applications. We analyzed the throughput of on-off channels with retransmissions, where we showed that even when packet sizes have small means and bounded support the variability in their sizes can greatly impact system performance. Specifically, if L (truncated at b) and A follow exponential distributions of rates μ and ν respectively with $\nu > \mu$, then as $b \to \infty$, the throughput vanishes to zero at a speed proportional to $e^{-(\nu-\mu)b}$. Next, we considered the distribution of the number of jobs in an $M/M/\infty$ queue with server failures. Here we showed that retransmissions can cause long-range dependence and quantified the impact of the maximum job sizes on the long-range dependence.

VII. Appendix

Since $\bar{F}(x)$ is a continuous function on $[0, \infty)$, we can define its inverse function $\bar{F}^{\leftarrow}(x) \triangleq \inf\{y : \bar{F}(y) < x\}$.

A. Proof of Proposition 3

First we prove the result for $N(b)$. Since the sequence $\{A_i\}$ is *i.i.d.*, upper bounding the packet size L by b yields $\mathbb{P}[N(b) > n] \le \mathbb{P}[A \le b]^n$.

Next, we study $T(b)$. Note that Condition 1 ensures the existence of γ^*, since the moment generating function $\int_0^\infty e^{\gamma s} d\mathbb{P}[A + U \le s | A \le b]$ is finite, continuous, and monotonically increasing for all $\gamma > 0$ [17]. For a random variable $\bar{N}$ that is independent of $\{A_i, U_i\}$ with $\mathbb{P}\left[\bar{N} > n\right] = \mathbb{P}[A \le b]^n$, $n = 0, 1, 2, \cdots$ (thus $\bar{N} \ge 1$), we have

$$T(b) \le_{st} \sum_{i=1}^{\bar{N}-1} (\bar{A}_i + U_i) + b, \tag{15}$$

where $\bar{A}_i, i \ge 1$ are *i.i.d.* and independent of $\bar{N}, \{U_i\}$ with $\mathbb{P}\left[\bar{A}_i > t\right] = \mathbb{P}[A > t | A \le b]$. Noting that the first term on the righthand side of (15) is a geometric sum of *i.i.d.* random variables, we derive the following defective renewal equation, for $t > 0$,

$$\mathbb{P}\left[\sum_{i=1}^{\bar{N}-1} (\bar{A}_i + U_i) > t\right] = \mathbb{P}\left[\bar{N} > 1, \bar{A}_1 + U_1 > t\right] + \mathbb{P}\left[\bar{N} > 1, \bar{A}_1 + U_1 \le t, \sum_{i=1}^{\bar{N}-1} (\bar{A}_i + U_i) > t\right]$$

$$= \mathbb{P}[A \le b]\mathbb{P}[\bar{A}_1 + U_1 > t] + \mathbb{P}[A \le b] \times \int_0^t \mathbb{P}\left[\sum_{i=2}^{\bar{N}-1} \bar{A}_i + U_i > t - u\right] d\mathbb{P}\left[\bar{A}_1 + U_1 \le u\right]. \tag{16}$$

Applying Theorem 7.1 of Chapter 5 in [18], we derive

$$\lim_{t\to\infty} e^{\gamma^*(t-b)}\mathbb{P}\left[\sum_{i=1}^{\bar{N}-1} (\bar{A}_i + U_i) + b > t\right] = \frac{\int_0^\infty e^{\gamma^* s}\mathbb{P}\left[A + U > s | A \le b\right] ds}{\int_0^\infty s e^{\gamma^* s} d\mathbb{P}\left[A + U \le s | A \le b\right]},$$

which, using $\int_0^\infty e^{\gamma^* s} d\mathbb{P}[A + U \le s | A \le b] = 1/\mathbb{P}[A \le b]$ to compute the numerator, completes the proof. ■

B. Proof of Theorem 1

Notice that the number of retransmissions is geometrically distributed given the unit size L,

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

and therefore,

$$\mathbb{P}[N(b) > n] = \mathbb{E}[(1 - \bar{G}(L))^n]. \tag{17}$$

I) First, we prove (4). Using Lemma 1 and Proposition 2, we obtain the upper bound

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[N(b) > n]}{\log n} \le \lim_{n\to\infty} \frac{\log \mathbb{P}[N(\infty) > n]}{\log n} = -\alpha,$$

which implies that, for $\epsilon > 0$, there exists n_0 such that

$$\inf_{n > n_0} \frac{\log \mathbb{P}\left[N(b) > n\right]}{-\alpha \log n} > 1 - \epsilon. \tag{18}$$

Next, we derive a lower bound. Condition 2 implies that for any $0 < \delta < 1/\alpha$, there exists x_δ, such that for all $x > x_\delta$, we have

$$\bar{F}(x)^{\frac{1}{\alpha}+\delta} \le \bar{G}(x) \le \bar{F}(x)^{\frac{1}{\alpha}-\delta}. \tag{19}$$

When $b > x_\delta$, the condition $n < \left(\bar{G}(b)\right)^{-\eta}$ implies $n <$

$\left(\bar{F}(b)\right)^{-\left(\frac{1}{\alpha}+\delta\right)\eta}$. Let $\zeta \triangleq (1/\alpha+\delta)\,\eta$, and we obtain $b > \bar{F}^{\leftarrow}(n^{-1/\zeta})$. Thus, there exists $b_0 > x_\delta$ such that for all $b > b_0$,

$$\begin{aligned}\mathbb{E}[(1-\bar{G}(L))^n] &= \mathbb{E}\left[(1-\bar{G}(L^*))^n|L^* \le b\right]\\ &\ge \frac{1}{\mathbb{P}[L^*\le b]}\mathbb{E}\left[\left(1-\bar{G}(L^*)\right)^n \mathbf{1}(x_\delta < L^* < b)\right]\\ &\ge \mathbb{E}\Big[\left(1-\bar{F}(L^*)^{\frac{1}{\alpha}-\delta}\right)^n \mathbf{1}\left(\bar{F}(b) < \bar{F}(L^*) < \bar{F}(x_\delta)\right)\Big]\\ &\ge \mathbb{E}\left[\left(1-\bar{F}(L^*)^{\frac{1}{\alpha}-\delta}\right)^n \mathbf{1}\left(\bar{F}(b) < \bar{F}(L^*)\right)\right]\\ &\quad - \left(1-\bar{F}(x_\delta)^{\frac{1}{\alpha}-\delta}\right)^n.\end{aligned}$$

Since $\bar{F}(x)$ is continuous, $\bar{F}(L^*)$ is a uniform random variable between 0 and 1, the preceding inequality implies

$$\begin{aligned}\mathbb{E}[(1-\bar{G}(L))^n] \ge \int_{\bar{F}(b)}^{1} &\left(1-u^{\frac{1}{\alpha}-\delta}\right)^n du\\ &- \left(1-\bar{F}(x_\delta)^{\frac{1}{\alpha}-\delta}\right)^n,\end{aligned}$$

which, letting $u^{1/\alpha-\delta} = v$ and noting $\bar{F}(b) < n^{-1/\zeta}$, yields

$$\begin{aligned}\mathbb{E}[(1-\bar{G}(L))^n] &\ge \int_0^1 (1-v)^n \,\frac{\alpha}{1-\alpha\delta} v^{\frac{\alpha}{1-\alpha\delta}-1} dv\\ &- \int_0^{n^{-1/\zeta}} \left(1-u^{\frac{1}{\alpha}-\delta}\right)^n du - \left(1-\bar{F}(x_\delta)^{\frac{1}{\alpha}-\delta}\right)^n\\ &\triangleq P_1 - P_2 - P_3.\end{aligned} \tag{20}$$

Recalling the property of Beta function, for fixed x,

$$B(x,y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt \sim \Gamma(x)y^{-x} \tag{21}$$

as $y\to\infty$, we obtain,

$$P_1 \sim \frac{\alpha\Gamma(\alpha_\delta)}{1-\alpha\delta} n^{-\alpha_\delta}, \quad \alpha_\delta = \frac{\alpha}{1-\alpha\delta}. \tag{22}$$

We know that

$$P_2 \le \frac{1}{n^{1/\zeta}}, \tag{23}$$

which, in conjunction with (20) and the fact that P_3 is exponentially bounded, implies that, for $0<\delta<1/\alpha-\zeta$, i.e., $1/\zeta > \alpha_\delta$, there exists $n_0 > 0$ such that for $n_0 < n < \left(\bar{G}(b)\right)^{-\eta}$,

$$\mathbb{E}[(1-\bar{G}(L))^n] \ge (1-\delta)\int_0^1 (1-v)^n \,\frac{\alpha}{1-\alpha\delta} v^{\frac{\alpha}{1-\alpha\delta}-1} dv.$$

Using (22) and the preceding lower bound, we obtain, for any fixed $\epsilon > 0$, there exists n_0 and b_0, where $\left(\bar{G}(b_0)\right)^{-\eta} > n_0$, such that for $b > b_0$,

$$\sup_{n_0<n<\left(\bar{G}(b)\right)^{-\eta}} \frac{\log\mathbb{P}\left[N\left(b\right)>n\right]}{-\alpha\log n} < 1+\epsilon. \tag{24}$$

Combining (18) and (24), we finish the proof of (4). In (23) we simply upper bound P_2 by $\bar{F}(b)$. Using a tighter bound we can prove the power law main body on a larger interval for n.

II) Next, we prove (5). Using $\mathbb{P}[L\le b^*]=1$ and (17) yields

$$\mathbb{P}[N(b)>n] \le (1-\bar{G}(b^*))^n,$$

which proves the upper bound.

Next, we prove the lower bound. For $\epsilon>0$, note that

$$\begin{aligned}\mathbb{P}[N(b)>n] &\ge \mathbb{E}\left[(1-\bar{G}(L)^n \mathbf{1}\left(b^*-\epsilon<L<b^*\right)\right]\\ &\ge \mathbb{P}[b^*-\epsilon<L<b^*](1-\bar{G}(b^*-\epsilon))^n.\end{aligned}$$

The definition of b^* implies $\mathbb{P}[b^*-\epsilon<L<b^*]>0$, and therefore,

$$\lim_{n\to\infty}\frac{\log\mathbb{P}\left[N\left(b\right)>n\right]}{n} \ge -\log\left(1-\bar{G}(b^*-\epsilon)\right),$$

which, by passing $\epsilon\to 0$ and using the continuity, finishes the proof of the lower bound. ■

C. Proof of Theorem 2

First, we prove equation (6). Since the proof is based on the same approach as in proving Theorem 3 in [8], here we only discuss the proof of the upper bound.

For any $\delta>0$, we have

$$\begin{aligned}\mathbb{P}[T(b)>t] &= \mathbb{P}\left[\sum_{i=1}^{N(b)-1}(U_i+A_i)+L>t\right]\\ &\le \mathbb{P}\left[\sum_{i=1}^{N(b)}(U_i+A_i)>t, N(b)\le\frac{t(1-\delta)}{\mathbb{E}[U+A]}\right]\\ &\quad+\mathbb{P}\left[N(b)>\frac{t(1-\delta)}{\mathbb{E}[U+A]}\right]+\mathbb{P}\left[L^*>t|L^*\le b\right]\\ &\triangleq I_1+I_2+I_3.\end{aligned} \tag{25}$$

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

$$I_1 \le \mathbb{P}\left[\sum_{i\le t(1-\delta)/\mathbb{E}[U+A]}(U_i+A_i)>t\right] = \mathbb{P}\left[\sum_{i\le\zeta t}X_i>\delta t\right],$$

which, by Chernoff bound, is bounded by $he^{-\zeta t}$ for $h,\zeta>0$.

Since I_3 is upper bounded by $2\mathbb{P}[L^*>t]$ that is independent of b for b large enough, the summation $I_1+I_2+I_3$ is dominated by I_2, which implies the upper bound for (6). The lower bound can be obtained by combining similar arguments as in the upper bound and the proof in [8].

Next, we prove (7). We begin with the upper bound. Recalling (1) and using Lemma 1, we upper bound the packet size L by b^* to obtain, almost surely,

$$T(b) \le \sum_{i=1}^{\bar{N}-1}(\bar{A}_i+U_i)+b^*, \tag{26}$$

where $\mathbb{P}\left[\bar{N}>n\right] = \left(1-\bar{G}(b^*)\right)^n$, $\mathbb{P}\left[\bar{A}_i>t\right] = \mathbb{P}[A>t|A\le b^*]$, $\{\bar{A}_i\}$ is a sequence of *i.i.d.* random variables independent from $\{U_i\}$ and $\bar{N}$ is independent of $\{\bar{A}_i, U_i\}$. Using the same approach as in computing (16), we obtain the upper bound.

Now, we prove the lower bound. For $0<\epsilon<b^*$, observe

$$\begin{aligned}\mathbb{P}[T(b)>t] &\ge \mathbb{P}\left[T(b)>t, b^*-\epsilon<L<b^*\right]\\ &\ge \mathbb{P}\left[b^*-\epsilon<L<b^*\right]\mathbb{P}\left[\sum_{i=1}^{N-1}(\underline{A}_i+U_i)+b^*-\epsilon>t\right],\end{aligned}$$

where $\mathbb{P}[\underline{N} > n] = \left(1 - \bar{G}(b^* - \epsilon)\right)^n$, $\mathbb{P}[\underline{A}_i > t] = \mathbb{P}[A > t | A \leq b^* - \epsilon]$, $\{\underline{A}_i\}$ is a sequence of *i.i.d.* random variables independent from $\{U_i\}$ and $\underline{N}$ is independent of $\{\underline{A}_i, U_i\}$. By the definition of b^*, we know $\mathbb{P}[b^* - \epsilon < L < b^*] > 0$. Using the same approach as in computing (16), we obtain

$$\lim_{t\to\infty} \frac{\log \mathbb{P}[T(b) > t]}{t} \geq -\gamma_\epsilon, \tag{27}$$

where γ_ϵ is the solution of

$$\int_0^\infty e^{\gamma s} d\mathbb{P}[A + U \leq s | A \leq b^* - \epsilon] = 1/\left(1 - \bar{G}(b^* - \epsilon)\right).$$

Passing $\epsilon \to 0$ and using the continuity, we finish the proof of the lower bound. ■

D. Proof of Theorem 3

First, we prove (10). Without loss of generality, we can assume that $\Phi(x)$ is absolutely continuous and strictly monotone since, by Proposition 1.5.8 of [15], one can always find an absolutely continuous and strictly monotone function

$$\Phi^*(x) = \alpha \int_{x_0}^x \Phi(s) s^{-1} ds, \quad x \geq x_0, \tag{28}$$

which, for large enough x_0, for all $x > x_0$, satisfies

$$\bar{F}^{-1}(x) \sim \Phi(\bar{G}^{-1}(x)) \sim \Phi^*(\bar{G}^{-1}(x)).$$

Therefore, as $n \to \infty$, we know $P(N(\infty) > n) \sim \Gamma(\alpha + 1)/\Phi(n)$ (c.f. [2]). Now, using this and Lemma 1, we have that, given any $0 < \epsilon < 1$, there exists an n_0 such that for all $n \geq n_0$,

$$\frac{P(N(b) > n)\Phi(n)}{\Gamma(\alpha+1)} < 1 + \epsilon.$$

Next, we prove the lower bound. Since $\Phi(x)$ is regularly varying function, there exists x_0 such that, the restriction of $\Phi(x)$ to $[x_0, \infty)$ has an inverse function $\Phi^{\leftarrow}(x)$. The condition (9) implies that, for $0 < \delta < 1$, there exists x_δ, such that for all $x > x_\delta$,

$$(1 - \delta)\bar{F}^{-1}(x) \leq \Phi(\bar{G}^{-1}(x)) \leq (1 + \delta)\bar{F}^{-1}(x),$$

and thus by choosing $b_0 > x_\delta > x_0$, we have that for all $b > b_0$,

$$\begin{aligned} \Phi^{\leftarrow}\left((1 - \delta)\bar{F}^{-1}(b)\right) &\leq \bar{G}^{-1}(b) \\ &\leq \Phi^{\leftarrow}\left((1 + \delta)\bar{F}^{-1}(b)\right). \end{aligned} \tag{29}$$

For $M > m > 0$ and $x_n > x_\delta$ with $\Phi^{\leftarrow}\left((1 - \delta)\bar{F}^{-1}(x_n)\right) = n/M$, we obtain, for $n \geq n_0$ and $b \geq b_0$,

$$\begin{aligned} &\mathbb{P}[N(b) > n] = \mathbb{E}\left[(1 - \bar{G}(L^*))^n | L^* \leq b\right] \\ &\geq \mathbb{E}\left[\left(1 - \bar{G}(L^*)\right)^n \mathbf{1}(x_n < L^* < b)\right] \\ &\geq \mathbb{E}\left[\left(1 - \frac{1}{\Phi^{\leftarrow}\left((1-\delta)\bar{F}^{-1}(L^*)\right)}\right)^n \mathbf{1}\left(\bar{F}(L^*) < \bar{F}(x_n)\right)\right] \\ &\quad - \mathbb{E}\left[\mathbf{1}\left(\bar{F}(L^*) \leq \bar{F}(b)\right)\right] \\ &\triangleq P_1 - P_2. \end{aligned}$$

For P_1, we obtain, by letting $z = n/\Phi^{\leftarrow}\left((1 - \delta)\bar{F}^{-1}(L^*)\right)$,

$$P_1\Phi(n) \geq \int_m^M \left(1 - \frac{z}{n}\right)^n (1 - \delta)\frac{\Phi(n)}{\Phi(\frac{n}{z})} \frac{\Phi'(\frac{n}{z})}{\Phi(\frac{n}{z})} \frac{n}{z^2} dz.$$

Since $\Phi(n)$ is a regularly varying function, by the Characterisation theorem of regular variation and the uniform convergence theorem of slowly varying functions, we obtain uniformly for any $0 < \delta < 1$ and $m \leq z \leq M$, $0 < m, M < \infty$, that there exists an n_1 such that for all $n > n_1$,

$$\frac{\Phi(n)}{\Phi(\frac{n}{z})} \geq (1 - \delta) z^\alpha.$$

Also, owing to (28), $\Phi'(n/z)/\Phi(n/z) = z\alpha/n$. In addition, there exists n_2 such that for all $n > n_2$, $(1 - z/n)^n > (1 - \delta)e^{-z}$. Hence, for all $n > n_0 = \max\{n_1, n_2\}$,

$$P_1\Phi(n) \geq \int_m^M (1 - \delta)^3 \alpha e^{-z} z^{\alpha - 1} dz.$$

The condition $n < \zeta/\bar{G}(b)$, in conjunction with (29), implies $\bar{F}(b) < (1 + \delta)/\Phi(n/\zeta)$. Therefore,

$$P_2\Phi(n) \leq \bar{F}(b)\Phi(n) \leq (1 + \delta)\frac{\Phi(n)}{\Phi(n/\zeta)}. \tag{30}$$

Since $\Phi(n)$ is regularly varying with $\alpha > 0$, for any given $\delta > 0$, we can choose n_0 large enough and ζ small enough such that for all $n > n_0$, $P_2\Phi(n) < \delta$. In (30) we simply upper bound P_2 by $\bar{F}(b)$, which is not tight. Using a tighter upper bound we can prove the power law main body on a larger interval for n.

Hence, using the fact $\int_0^\infty \alpha e^{-z} z^{\alpha-1} dz = \Gamma(\alpha + 1)$, for any $0 < \epsilon < 1$, we can choose M large enough and m, ζ, δ small enough, such that for $b > b_0$ and $n_0 < n < \zeta/\bar{G}(b)$,

$$\frac{P(N(b) > n)\Phi(n)}{\Gamma(\alpha+1)} > 1 - \epsilon.$$

The proof of (11) follows from (25) and (10). ■

E. Proof of Lemma 2

We first show that, for any fixed $\eta_2 > \eta_1 > 1$ and $0 < \epsilon < \eta_1$, there exists b_0 such that for all $b > b_0$, and for $\left(\bar{G}(b)\right)^{-\eta_1} < n \left(\bar{G}(b)\right)^{-\eta_2}$,

$$-n^{1 - \frac{1}{\eta_2 + \epsilon}} \leq \log \mathbb{P}[N(b) > n] \leq -n^{1 - \frac{1}{\eta_1 - \epsilon}} \tag{31}$$

Condition 2 implies that for any $0 < \delta < 1/\alpha$, there exists x_δ, such that for all $x > x_\delta$, we have

$$\bar{F}(x)^{\frac{1}{\alpha} + \delta} \leq \bar{G}(x) \leq \bar{F}(x)^{\frac{1}{\alpha} - \delta}. \tag{32}$$

Since $\bar{F}(x)$ is eventually non-increasing, by equation (19), the condition $\left(\bar{G}(b)\right)^{-\eta_1} < n < \left(\bar{G}(b)\right)^{-\eta_2}$ implies that

$$\bar{F}^{\leftarrow}\left(n^{-1/\zeta_2}\right) < b < \bar{F}^{\leftarrow}\left(n^{-1/\zeta_1}\right) \tag{33}$$

where $\zeta_2 \triangleq (1/\alpha + \delta)\eta_2$ and $\zeta_1 \triangleq (1/\alpha - \delta)\eta_1$. Choosing $0 < \delta < (\eta_1 - 1)/(\alpha\eta_1)$, we obtain $\zeta_2 > \zeta_1 > 1/\alpha$. Let $V = \bar{F}(L^*)$.

For the upper bound, by Lemma 1 and (33), we obtain,

$$\begin{aligned}\mathbb{P}[N(b)>n] &= \mathbb{E}\left[(1-\bar{G}(L^*))^n | L^* \le b\right] \\ &\le \frac{1}{1-n^{-1/\zeta_2}}\mathbb{E}\Big[\left(1-\bar{G}(L^*)\right)^n \\ &\qquad \mathbf{1}\left(x_\delta < L^* \le \bar{F}^{\leftarrow}\left(n^{-1/\zeta_1}\right)\right)\Big] \\ &\quad + \frac{1}{1-n^{-1/\zeta_2}}(1-\bar{G}(x_\delta))^n.\end{aligned}$$

Therefore, there exists n_0, such that for all $n > n_0$, $\mathbb{P}[N(b) > n]$ is upper bounded by

$$\begin{aligned}&(1+\delta)\mathbb{E}\left[\left(1-V^{\frac{1}{\alpha}+\delta}\right)^n \mathbf{1}\left(\frac{1}{n^{1/\zeta_1}} < V < \bar{F}(x_\delta)\right)\right] \\ &\quad + (1+\delta)(1-\bar{G}(x_\delta))^n \\ &\le (1+\delta)\left(1-\frac{1}{n^{1/(\zeta_1\alpha)+\delta/\zeta_1}}\right)^n + (1+\delta)(1-\bar{G}(x_\delta))^n \\ &\le (1+2\delta)e^{-n^{1-1/(\zeta_1\alpha)-\delta/\zeta_1}},\end{aligned}$$

which implies, there exists b_0 large enough, such that for any given $\eta_1 > 1$, and $0 < \epsilon < \eta_1$, for all $b > b_0$ and $n > \left(\bar{G}(b)\right)^{-\eta_1}$,

$$\log \mathbb{P}\left[N(b) > n\right] \le -n^{1-\frac{1}{\alpha(\zeta_1-\epsilon)}}. \tag{34}$$

Now, we can prove Equation (13) in a similar way to proving Equation (6) by appealing to Equation (25). To prove the lower bound, for $0 < \epsilon < 1/\alpha$, we obtain

$$\begin{aligned}&\mathbb{P}[N(b)>n] = \mathbb{E}\left[(1-\bar{G}(L^*))^n | L^* \le b\right] \\ &\ge \mathbb{E}\Big[\left(1-\bar{G}(L^*)\right)^n \\ &\qquad \mathbf{1}\left(\bar{F}^{\leftarrow}\left(n^{-1/\zeta_2}\right) < L^* < \bar{F}^{\leftarrow}\left(n^{-1/\zeta_1}\right)\right)\Big] \\ &\ge \mathbb{E}\left[\left(1-V^{\frac{1}{\alpha}-\epsilon}\right)^n \mathbf{1}\left(\frac{1}{n^{1/\zeta_1}} < V < \frac{1}{n^{1/\zeta_2}}\right)\right] \\ &\ge \left(\frac{1}{n^{1/\zeta_2}} - \frac{1}{n^{1/\zeta_1}}\right)\left(1-\frac{1}{n^{1/(\alpha\zeta_2)-\epsilon/\zeta_2}}\right)^n \\ &\sim \frac{1}{n^{1/\zeta_2}}e^{-n^{1-1/(\alpha\zeta_2)-\epsilon/\zeta_2}},\end{aligned}$$

which implies, for $n > n_0$ with n_0 large enough,

$$\log \mathbb{P}\left[N(b) > n\right] \ge -n^{1-\frac{1}{\alpha(\zeta_2+\epsilon)}}. \tag{35}$$

Combining (34), (35), we finish the proof of (12). We obtain (13) in a similar way as in the proof of Theorem 2. ■

F. Proof of Theorem 4

Since $\mathbb{E}[e^{\theta A}] < \infty$ for some $\theta > 0$, we can always find a random variable X and $t_0, \delta > 0$ with $A \overset{d}{\le} X$ and

$$\mathbb{P}[X > t] = \begin{cases} 1 & \text{if } t \le t_0, \\ e^{-\delta(t-t_0)} & \text{if } t > t_0. \end{cases} \tag{36}$$

It can be checked that, $\mathbb{P}[X > t+s] \le \mathbb{P}[X > t]\mathbb{P}[X > s]$ for any $t, s \ge 0$. Therefore, for $t \ge 0$ and any positive random variable Y that is independent of X with $\mathbb{P}[X > Y] > 0$, we obtain

$$\mathbb{P}\left[X > Y + t | X > Y\right] \le \mathbb{P}[X > t]. \tag{37}$$

Denote by $\Psi(i), i \ge 1$ the index of the channel available period within which L_i succeeds in transmission. Immediately after the successful transmission of the packet L_i, we denote by $\{\tau_i\}$ the remaining time and by $\{\sigma_i\}$ the elapsed time, respectively, within the available period $A_{\Psi(i)}$.

First, we prove the upper bound using the coupling argument. Using (36), we can construct in the same probability space an *i.i.d.* sequence $\{X_i^{(1)}\}_{i\ge1}$ with $X_i^{(1)} \ge A_i$ and $X_i^{(1)} \overset{d}{=} X$ for all $i \ge 1$. Define $\bar{\tau}_i = X_{\Psi(i)}^{(1)} - \sigma_i$; obviously, $\bar{\tau}_i \ge \tau_i$. For a successful transmission with $L_i = \sigma_i$, we obtain, by (37),

$$\begin{aligned}\mathbb{P}[\tau_i > t | L_i = \sigma_i] &\le \mathbb{P}[\bar{\tau}_i > t | L_i = \sigma_i] \\ &= \mathbb{P}[X_1^{(1)} > L_i + t | X_1^{(1)} > L_i] \le \mathbb{P}[X > t]. \end{aligned} \tag{38}$$

Now, we will construct a new system where we always have $L_i = \sigma_i$, and thus $\mathbb{P}[\tau_i > t | L_i = \sigma_i] = \mathbb{P}[\tau_i > t]$. We continue to use $\Psi(i), i \ge 1$ as the index of the channel available period within which L_i succeeds in transmission in the newly constructed system. At the beginning of A_1, we replace A_1 by $X_1^{(1)}$ and start transmitting L_1. Then, immediately after each successful transmission, say, for packet L_i in the available period $A_{\Psi(i)}$, in view of (38), we can construct a new available period by replacing τ_i with $X_i^{(2)}$, where $X_i^{(2)} \overset{d}{=} X$ and $X_i^{(2)} \ge \tau_i$. Note that in this construction we change $A_{\Psi(i)}$ for all i and other available periods are the same as the original ones. Then, let the system continue its operation by following this construction. Noting that X has a constant hazard rate δ if $X \ge t_0$, the random variable $X_i^{(2)}$ is independent of all the random variables that appear before $X_i^{(2)}$ is generated in this new system. Denote by $\underline{T}_i, i \ge 1$ the transmission duration for packet L_i in this new system excluding the time that this packet spends in the constructed interval $X_i^{(2)}$ and the unavailable period $U_{\Psi(i)}$ that immediately follows $X_i^{(2)}$; clearly $\sum_{i=1}^n \underline{T}_i \le \sum_{i=1}^n T_i$ for all n. In addition, this construction implies that $\{\underline{T}_i\}_{i\ge1}$ is an *i.i.d.* sequence. If $X_1^{(1)} > L_1$, the first transmission in the available period $X_1^{(1)}$ (replacing A_1) succeeds, and thus $\underline{T}_1 = 0$. If $X_1^{(1)} \le L_1$, since the first transmission fails and we need to wait until the beginning of the second available period A_2 to retransmit L_1, $\underline{T}_1 \overset{d}{=} T_1$, where T_1 is the transmission time of the first packet in the original system. Hence, we obtain $\underline{T}_i \overset{d}{=} T_1 \mathbf{1}(L_1 > X)$, where X is independent of L_1 and $\{A_i, U_i\}_{i\ge1}$.

Therefore, as $n \to \infty$, we obtain, using the law of large numbers,

$$\Lambda_n(b) \le \frac{\sum_{i=1}^n L_i}{\sum_{i=1}^n \underline{T}_i} = \frac{\sum_{i=1}^n L_i}{n}\frac{n}{\sum_{i=1}^n \underline{T}_i} \to \frac{\mathbb{E}[L]}{\mathbb{E}[\underline{T}_1]}. \tag{39}$$

Now, we need to compute, for $t > 0$,

$$\mathbb{P}[\underline{T}_1 > t] = \mathbb{P}[T_1 > t, L_1 > X].$$

Using the same x_ϵ as in (19), we obtain, by the independence of L_1 and X,

$$\mathbb{P}[T_1 > t, L_1 > X] \ge \mathbb{P}[T_1 > t, L_1 > X, L_1 > x_\epsilon]$$

So, $\mathbb{P}[T_1 > t, L_1 > X] \geq \mathbb{P}[T_1 > t, L_1 > x_\epsilon, X < x_\epsilon]$

$$= \mathbb{P}[T_1 > t, L_1 > x_\epsilon]\mathbb{P}[X < x_\epsilon], \tag{40}$$

where, choosing x_ϵ large enough, we can always make $\mathbb{P}[X < x_\epsilon] > 0$. From the proof of Theorem 2, we know that, for any $0 < \epsilon < 1$, there exist t_0 and b_0 such that for $b > b_0$ and $t_0 < t < \bar{G}(b)^{-(1-\epsilon)}$,

$$\mathbb{P}[T_1 > t, L_1 > x_\epsilon] > \frac{1}{t^{(1+\epsilon)\alpha}},$$

which, by (40), implies,

$$\log \mathbb{E}[\underline{T}_1] \geq \log \int_{t_0}^{\bar{G}(b)^{-(1-\epsilon)}} \frac{\mathbb{P}[X < x_\epsilon]}{t^{(1+\epsilon)\alpha}} dt$$
$$\sim -(1-\epsilon)(1-(1+\epsilon)\alpha)\log \bar{G}(b), \tag{41}$$

as $b \to \infty$. Using (41), (39), and passing $\epsilon \to 0$, we prove, as $b \to \infty$,

$$\limsup_{n\to\infty} \frac{\log \Lambda_n(b)}{\log \bar{G}(b)} \lesssim 1 - \alpha.$$

Next, we prove the lower bound. In each available period $A_{\Psi(i)}, i \geq 1$ that contains a successful transmission, we postpone the transmission of the new packet L_{i+1} until the beginning of the next available period $A_{\Psi(i)+1}$. It is easy to see that, this construction increases the total transmission time, and also the durations for transmitting $L_i, i \geq 1$ are *i.i.d.* random variables. Thus, the law of large numbers can be applied. Based on similar arguments in deriving the upper bound, we can prove, as $b \to \infty$,

$$\liminf_{n\to\infty} \frac{\log \Lambda_n(b)}{\log \bar{G}(b)} \gtrsim 1 - \alpha.$$

Combining the lower and upper bounds, we complete the proof. ■

G. Proof of Theorem 6

Let $T_i(b)$ be the processing time for job i. For $\alpha = \mu/\nu$ and $\bar{G}(b) = e^{-\nu b}$, we obtain, by Theorem 2, that for any $0 < \epsilon < 1$, there exist t_0 and b_0 such that for any $b > b_0$ and $t_0 < t < \left(\bar{G}(b)\right)^{-(1-\epsilon)}$,

$$\mathbb{P}[T_1(b) > t] > \frac{1}{t^{(1+\epsilon)\alpha}}.$$

Using Theorem 5,

$$\int_0^\infty \mathrm{Cov}(M(t), M(t+s)) = \lambda \int_0^\infty \int_s^\infty \mathbb{P}[T_1(b) > t] dt\ ds$$
$$\geq \lambda \int_{t_0}^{\left(\bar{G}(b)\right)^{-(1-\epsilon)}} \int_s^{\left(\bar{G}(b)\right)^{-(1-\epsilon)}} \frac{1}{t^{(1+\epsilon)\alpha}} dt\ ds$$
$$= \frac{\lambda}{(1+\epsilon)\alpha - 1}\Bigg[\frac{\left(\bar{G}(b)\right)^{-(1-\epsilon)(2-(1+\epsilon)\alpha)}}{(2-(1+\epsilon)\alpha)} - \frac{t_0^{(2-(1+\epsilon)\alpha)}}{(2-(1+\epsilon)\alpha)}$$
$$- \left(\bar{G}(b)\right)^{-(1-\epsilon)(2-(1+\epsilon)\alpha)} + t_0\left(\bar{G}(b)\right)^{-(1-\epsilon)(1-(1+\epsilon)\alpha)}\Bigg]$$
$$= \frac{\lambda e^{(1-\epsilon)(2\nu-(1+\epsilon)\mu)b}}{(2-(1+\epsilon)\alpha)}\Bigg(1 - \frac{t_0^{(2-(1+\epsilon)\alpha)}}{(1+\epsilon)\alpha - 1} e^{-(1-\epsilon)(2\nu-(1+\epsilon)\mu)b}$$
$$+ t_0 \frac{2-(1+\epsilon)\alpha}{(1+\epsilon)\alpha - 1} e^{-(1-\epsilon)\nu b}\Bigg)$$
$$= \frac{\lambda e^{(1-\epsilon)(2\nu-(1+\epsilon)\mu)b}}{(2-(1+\epsilon)\alpha)}(1+o(1)).$$

Passing $b \to \infty$ and then $\epsilon \to 0$ yields

$$\log \int_0^\infty \mathrm{Cov}(M(t), M(t+s)) ds \gtrsim (2\nu - \mu)b,$$

which shows the lower bound.
Next we prove the upper bound. For $\eta_1 > 1$, we obtain

$$\int_0^\infty \int_s^\infty \mathbb{P}[T_1(b) > t] dt\ ds$$
$$= \int_0^{\left(\bar{G}(b)\right)^{-\eta_1}} \int_s^\infty \mathbb{P}[T_1(b) > t] dt\ ds$$
$$+ \int_{\left(\bar{G}(b)\right)^{-\eta_1}}^\infty \int_s^\infty \mathbb{P}[T_1(b) > t] dt\ ds$$
$$\triangleq I_1 + I_2.$$

By Lemma 2, we have that, for a fixed $\epsilon > 0$ and $\eta_1 = 1 + 2\epsilon$ there exists b_0 such that for all $b > b_0$ and $t > \left(\bar{G}(b)\right)^{-\eta_1}$,

$$\mathbb{P}[T_1(b) > t] \leq \exp\left(-t^{1-\frac{1}{\eta_1 - \epsilon}}\right).$$

For $0 < \zeta = 1 - 1/(\eta_1 - \epsilon) < 1$, using the above inequality in I_2, we have that for all $b > b_0$,

$$I_2 \leq \int_{\left(\bar{G}(b)\right)^{-\eta_1}}^\infty \int_s^\infty e^{-t^\zeta} dt\ ds$$
$$\leq \int_{\left(\bar{G}(b)\right)^{-\eta_1}}^\infty \left(\frac{1}{\zeta} - 1\right)\Gamma\left(\frac{1}{\zeta}, s^\zeta\right) ds,$$

where $\Gamma(x, s) = \int_s^\infty t^{x-1}e^{-t} dt$ is the upper incomplete Gamma function. For $\epsilon > 0$, there exists s_0 such that for all $s > s_0$, $\Gamma(x, s) \leq (1+\epsilon)s^{x-1}e^{-s}$. Now, we can choose b_0 such that for all $b > b_0$, $\left(\bar{G}(b)\right)^{-\eta_1} > s_0$. Hence, as $b \to \infty$,

$$I_2 \leq \int_{\left(\bar{G}(b)\right)^{-\eta_1 \zeta}}^\infty (1-\epsilon)\left(\frac{1}{\zeta} - 1\right)^2 s^{\frac{2}{\zeta} - 1} e^{-s} ds \to 0.$$

To obtain an upper bound on I_1, we use Lemma 1. Using Proposition 2, there exists $t_0 > 0$ such that for all $t > t_0$,

$$P(T(\infty) > t) < \frac{1}{t^{\alpha - \epsilon}}.$$

Now, we can choose b_0 such that for all $b > b_0$, $\left(\bar{G}(b)\right)^{-\eta_1} > t_0$, and thus,

$$I_1 \leq \int_0^{t_0} \int_s^\infty \mathbb{P}[T_1(b) > t] dt\ ds$$
$$+ \int_{t_0}^{\left(\bar{G}(b)\right)^{-\eta_1}} \int_s^\infty \mathbb{P}[T_1(\infty) > t] dt\ ds$$
$$\triangleq I_{11} + I_{12}.$$

For I_{11},

$$I_{11} = \int_0^{t_0}\int_s^{t_0} \mathbb{P}[T_1(b) > t]dt\ ds + \int_0^{t_0}\int_{t_0}^{\infty} \mathbb{P}[T_1(b) > t]dt\ ds \leq t_0^2 + \int_0^{t_0}\int_{t_0}^{\infty} \frac{1}{t^{\alpha-\epsilon}} dt\ ds,$$

which, due to $1 < \alpha < 2$, is upper bounded by a finite constant independent of b.

For I_{12}, we have that for all $b > b_0$,

$$I_{12} \leq \int_{t_0}^{(\bar{G}(b))^{-\eta_1}} \int_s^{\infty} \frac{1}{t^{\alpha-\epsilon}} dt\ ds \leq \frac{1}{(\alpha - \epsilon - 1)(2 + \epsilon - \alpha)} e^{\eta_1 b((2+\epsilon)\nu - \mu)}.$$

Now, by letting $\epsilon \to 0$, and $\eta_1 \to 1$ we have that

$$\log \int_0^{\infty} \mathrm{Cov}(M(t), M(t+s))ds \lesssim (2\nu - \mu)b.$$

■

References

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

[2] P. R. Jelenković and J. Tan, "Characterizing heavy-tailed distributions induced by retransmissions," *Advances in Applied Probability*, vol. 45, no. 1, 2013, to appear, Eprint arXiv: 0709.1138v2.

[3] ——, "Is ALOHA causing power law delays?" in *Proceedings of the 20th International Teletraffic Congress,* Ottawa, Canada, ser. Lecture Notes in Computer Science, vol. 4516. Springer-Verlag, June 2007, pp. 1149–1160.

[4] S. Asmussen, P. Fiorini, L. Lipsky, T. Rolski, and R. Sheahan, "Asymptotic behavior of total times for jobs that must start over if a failure occurs," *Mathematics of Operations Research*, vol. 33, no. 4, pp. 932–944, November 2008.

[5] J. Nair, M. Andreasson, L. L. H. Andrew, S. H. Low, and J. Doyle, "File fragmentation over an unreliable channel," in *INFOCOM*, 2010, pp. 965–973.

[6] J. Tan, Y. Yang, N. B. Shroff, and H. E. Gamal, "Delay asymptotics with retransmissions and fixed rate codes over erasure channels," in *INFOCOM*, 2011, pp. 1260–1268.

[7] J. Tan, W. Wei, B. Jiang, N. B. Shroff, and D. F. Towsley, "Can multipath mitigate power law delays?: Effects of parallelism on tail performance," in *SIGMETRICS*, 2010, pp. 381–382.

[8] P. R. Jelenković and J. Tan, "Can retransmissions of superexponential documents cause subexponential delays?" in *Proceedings of IEEE INFOCOM'07*, Anchorage, Alaska, USA, 2007, pp. 892–900.

[9] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović, "Power law and exponential decay of inter contact times between mobile devices," in *MobiCom '07: Proceedings of the 13th annual ACM international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2007, pp. 183–194.

[10] W. E. Leland, W. Willinger, M. S. Taqqu, and D. V. Wilson, "On the self-similar nature of Ethernet traffic," *ACM SIGCOMM Computer Communication Review*, vol. 25, no. 1, pp. 202–213, 1995.

[11] P. Jacquet, A. M. Naimi, and G. Rodolakis, "Asymptotic delay analysis for cross-layer delay-based routing in ad hoc networks," *Advances in Multimedia*, vol. 2007, p. 11, 2007, article ID 90879, doi:10.1155/2007/90879.

[12] N. A. Ali, E. Ekram, A. Eljasmy, and K. Shuaib, "Measured delay distribution in a wireless mesh network test-bed," *Computer Systems and Applications, ACS/IEEE International Conference on*, pp. 236–240, 2008.

[13] W.-B. Gong, Y. Liu, V. Misra, and D. Towsley, "Self-similarity and long range dependence on the Internet: a second look at the evidence, origins and implications," *Computer Networks*, vol. 48, pp. 377–399, June 2005.

[14] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *Proceedings of ACM E-WIND '05*, Philadelphia, Pennsylvania, USA, 2005, pp. 5–10.

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

[16] M. Parulekar and A. Makowski, "$M/G/\infty$ input processes: a versatile class of models for network traffic," in *INFOCOM '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 2, April 1997, pp. 419–426.

[17] K. L. Chung, *A course in probability theory*, 2nd ed. Academic Press, New York, 1974.

[18] S. Asmussen, *Applied Probability and Queues*, 2nd ed. Springer-Verlag, 1987.

Jian Tan is a Research Staff Member in IBM T.J. Watson Research Center, New York. He was a postdoctoral researcher at the Networking and Communications Research Lab in The Ohio State University from 2009 to 2010. He received B.S. from University of Science and Technology of China in 2002, and M.S. and Ph.D. degrees from Columbia University in 2004 and 2009, all in Electrical Engineering. He interned at Lucent Bell Laboratories, Murray Hill, NJ, during the summers of 2005 and 2006, and at Microsoft Research, Cambridge, UK, in the winter of 2007. His current research interests focus on big data platforms, cluster computing and related applications.

B. T. Swapna is a doctoral student in the Department of Electrical and Computer Engineering at The Ohio State University. She received her B.Tech and M.Tech degrees from the Indian Institute of Technology, Madras in 2008, both in Electrical Engineering.

Her research interests are in the areas of communication networks, social and information networks.

Ness B. Shroff (S'91-M'93-SM'01-F'07) received his Ph.D. degree from Columbia University, NY in 1994 and joined Purdue university immediately thereafter as an Assistant Professor. At Purdue, he became Professor of the school of Electrical and Computer Engineering in 2003 and director of CWSA in 2004, a university-wide center on wireless systems and applications. In July 2007, he joined the ECE and CSE departments at The Ohio State University, where he holds the Ohio Eminent Scholar Chaired Professorship of Networking and Communications.

His research interests span the areas of wireless and wireline communication networks. He is especially interested in fundamental problems in the design, performance, control, and security of these networks. Dr. Shroff is a past editor for IEEE/ACM Trans. on Networking and the IEEE Communications Letters. He currently serves on the editorial board of the Computer Networks Journal, IEEE Network Magazine and the Network Science journal. He has served on the technical and executive committees of several major conferences and workshops.

Dr. Shroff is a Fellow of the IEEE, and a National Science Foundation CAREER awardee. His papers have received numerous awards at top-tier conferences and journals, including the IEEE WiOPT best student paper award, the IEEE INFOCOM 2008 best paper award, the IEEE INFOCOM 2006 best paper award, the IEEE IWQoS 2006 best student paper award, the 2005 best paper of the year award for the Journal of Commnications and Networking, the 2003 best paper of the year award for Computer Networks (his INFOCOM 2005 paper was also selected as one of two runner-up papers for the best paper award).