Is ALOHA Causing Power Law Delays?

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

Abstract. Renewed interest in ALOHA-based Medium Access Control (MAC) protocols stems from their proposed applications to wireless ad hoc and sensor networks that require distributed and low complexity channel access algorithms. In this paper, unlike in the traditional work that focused on mean value (throughput) and stability analysis, we study the distributional properties of packet transmission delays over an ALOHA channel. We discover a new phenomenon showing 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. From an engineering perspective, our results imply that the variability of packet sizes and number of active users need to be taken into consideration when designing robust MAC protocols, especially for ad-hoc/sensor networks where other factors, such as link failures and mobility, might further compound the problem.

Key words: ALOHA, medium access control, power laws, heavy-tailed distributions, light-tailed distributions, ad-hoc/sensor networks.

1 Introduction

ALOHA represents one of the first and most basic distributed Medium Access Control (MAC) protocols [1]. It is easy to implement since it does not require any user coordination or complicated controls and, thus, represents a basis for many modern MAC protocols, e.g., Carrier Sense Multiple Access (CSMA). Basically, ALOHA enables multiple users to share a common communication medium (channel) in a completely uncoordinated manner. Namely, a user attempts to send a packet over the common channel and, if there are no other user (packet) transmissions during the same time, the packet is considered successfully transmitted. Otherwise, if the transmissions of more than one packet (user) overlap, we say that there is a collision and the colliding packets need to

This work is supported by the National Science Foundation under grant number CNS-0435168.
Published in Proceedings of ITC-20, Ottawa, Canada, June 17-21, 2007.

be retransmitted. Each user retransmits a packet after waiting for an independent (usually exponential/geometric) period of time, making ALOHA entirely decentralized and asynchronous. The desirable properties of ALOHA, including its low complexity and distributed/asynchronous nature, make it especially beneficial for wireless sensor networks with limited resources as well as for wireless ad hoc networks that have difficulty in carrier sensing due to hidden terminal problems and mobility. This explains the recent renewed interests in ALOHA type protocols.

Traditionally, the performance evaluation of ALOHA has focused on mean value (throughput) and stability analysis, the examples of which can be found in every standard textbook on networking, e.g., see [3, 9, 8]; for more recent references see [7] and the references therein (due to space limitations, we do not provide comprehensive literature review on ALOHA in this paper). However, it appears that there are no explicit and general studies (more than two users) of the distributional properties of ALOHA, e.g., delay distributions. In this regard, in Subsection 2.1, we consider a standard finite population ALOHA model with variable length packets [4, 2] that have an asymptotically exponential tail. Surprisingly, we discover a new phenomenon that the distribution of the number of retransmissions (collisions) and time between two successful transmissions follow power law distributions, as stated in Theorem 1. Informally, our theorem shows that when the exponential decay rate of the packet distribution is smaller than the parameter of the exponential backoff distribution, even the finite population ALOHA has zero throughput. This is contrary to the common belief that the finite population ALOHA system always has a positive, albeit possibly small, throughput. Furthermore, even when the long term throughput is positive, the high variability of power laws (infinite variance when the power law exponent is less than 2) may cause long periods of very high congestion/low throughput. It also may appear counterintuitive that the system is characterized by power laws even though the distributions of all the variables (arrivals, backoffs and packets) of the system are of exponential type. However, this is in line with the very recent results in [5, 10, 6], which show that job completion times in systems with failures where jobs restart from the very beginning exhibit similar power law behavior. Our study in [6] was done in the communication context where job completion times are represented by document/packet transmission delays. It may also be worth noting that [6] reveals the existence of power law delays regardless of how light or heavy the packet/document and link failure distributions may be (e.g., Gaussian), as long as they have proportional hazard functions. Furthermore, from a mathematical perspective, our Theorem 1 analyzes a more complex setting than the one in [6, 10] and, thus, requires a novel proof. Hence, when compared with [6, 10], this paper both discovers a new related phenomenon in a communication MAC layer application area and provides a novel analysis of it.

As already stated in the abstract, the preceding power law phenomenon is a result of combined effects of packet variability and collisions. Hence, one can see easily that the power law delays can be eliminated by reducing the variability of packets. Indeed, for slotted ALOHA with constant size packets the delays are geometrically distributed. However, we show in Section 3 that, when the number of users sharing the channel is geometrically distributed, the slotted ALOHA exhibits power law delays as well.

In Section 4, we illustrate our results with simulation experiments, which show that the asymptotic power law regime is valid for relatively small delays and reasonably large probability values. Furthermore, the distribution of packets in practice might have a bounded support. To this end, we show by a simulation experiment that this situation results in distributions that have power law main body with an exponentiated (stretched) support in relation to the support of the packet size/number of active users. Hence, although exponentially bounded, the delays may be prohibitively long.

In practical applications, we may have combined effects of both variable packets and a random number of users, implying that the delay and congestion is likely to be even worse than predicted by our results. Thus, from an engineering perspective, one has to pay special attention to the packet variability and the number of users when designing robust MAC protocols, especially for ad-hoc/sensor networks where link failures [6], mobility and many other factors might further worsen the performance.

2 Power Laws in the Finite Population ALOHA with Variable Size Packets

In this section we show that the variability of packet sizes, when coupled with the contention nature of ALOHA, is a cause of power law delays. This study is motivated by the well-known fact that packets in todays Internet have variable sizes. To further emphasize that packet variability is a sole cause of power laws, we assume a finite population ALOHA model where each user can hold (queue) up to one packet at the time since the increased queueing only further exacerbates the problem. In addition, in Section 3 we show that the user variability in an infinite population model may be a cause of power law delays as well. In the remainder of this section, we describe the model and introduce the necessary notation in Subsection 2.1 and then in Subsection 2.2 we formulate and prove our main result on the logarithmic asymptotics of the transmission delay in Theorem 1.

2.1 Model Description

Consider $M \geq 2$ users sharing a common communication link (channel) of unit capacity. Each user can hold at most one packet in its queue and, when the queue is empty, a new packet is generated after an independent (from all other variables) exponential time with mean $1/\lambda$. Each packet has an independent length that is equal in distribution to a generic random variable L. A user with a newly generated packet attempts its transmission immediately and, if there are no other users transmitting during the same time, the packet is considered successfully transmitted. Otherwise, if the transmissions of more than one packet overlap, we say that there is a collision and the colliding packets need to be retransmitted; for a visual representation of the system see Figure 1. After a collision, each participating user waits (backoffs) for an independent exponential period of time with mean $1/\nu$ and then attempts to retransmit its packet. Each such user continues this procedure until its packet is successfully transmitted and then it generates a new packet after an independent exponential time of mean $1/\lambda$. Let $\{U(t)\}_{t\geq 0}$ denote the number of users that are in backoff state at time t.

Without loss of generality, assume that there is a successful transmission at time $t = 0$ and let $\{T_i\}_{i\geq 1}$ be an increasing sequence of positive time points when either a collision or successful transmission occurs. Let N be the smallest index i such that at time $T \equiv T_N$ there was a successful transmission. We will study the asymptotic properties of the distributions of N and T, representing the total number of transmission attempts per one successful transmission and the time between the two consecutive successful transmissions, respectively.

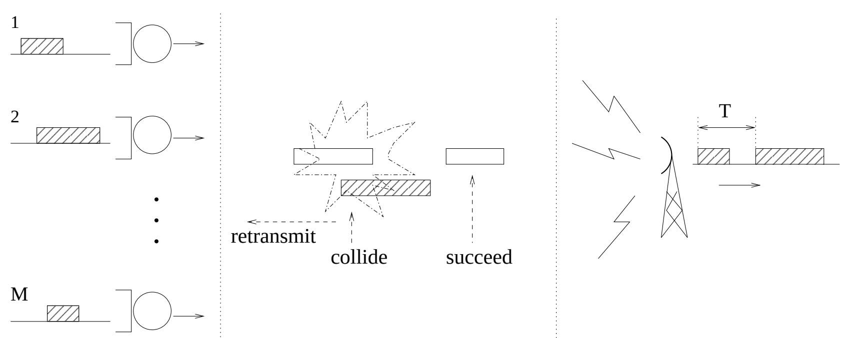

Fig. 1. Finite population ALOHA model with variable packet sizes.

2.2 Power Law Asymptotics

The following theorem on the logarithmic asymptotics of the number of transmission attempts N per successful transmission and delay T between two successful transmissions is our main result of this section.

Theorem 1. *If*

$$\lim_{x\to\infty} \frac{\log \mathbb{P}[L > x]}{x} = -\mu, \quad \mu > 0, \tag{1}$$

then, we have

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[N > n]}{\log n} = -\frac{M\mu}{(M-1)\nu} \tag{2}$$

and

$$\lim_{t\to\infty} \frac{\log \mathbb{P}[T>t]}{\log t} = -\frac{M\mu}{(M-1)\nu}. \tag{3}$$

Remark 1. The proof of this result reveals that if the transmission delay is long, then the shortest packet from all M users will be the most likely one that is successfully transmitted, i.e., ALOHA is unfair to longer packets.

Remark 2. This theorem indicates that the distribution tails of N and T are essentially power laws when the packet distribution is approximately exponential ($\approx e^{-\mu x}$). Thus, the finite population ALOHA may exhibit high variations and possible zero throughput. More precisely, by the strong law of large numbers for stationary and ergodic point processes, the system has zero throughput when $0 < M\mu/(M-1)\nu < 1$; and when $1 < M\mu/(M-1)\nu < 2$, the transmission time has finite mean but infinite variance. Furthermore, for large M, $M\mu/(M-1) \approx \mu/\nu$ and thus, the system has zero throughput if the backoff parameter $\nu \gtrapprox \mu$. It might be worth noting that this may even occur when the expected packet length is much smaller than the expected backoff time $\mathbb{E}L \ll 1/\nu$.

Proof. Let us first prove equation (2) assuming that at time $t=0$ a collision happens and all users have a packet waiting to be send, i.e., $U(0)=M$. Since each user i has an equal probability $1/M$ of being the first one to attempt a transmission, and $e^{-L_i(M-1)\nu}$ is the conditional probability, given L_i, that such an attempt is successful, we obtain, for $x_\epsilon > 0$,

$$\begin{aligned}\mathbb{P}[N>n] &= \mathbb{E}\left[\left(1-\frac{1}{M}\left(\sum_{i=1}^{M} e^{-L_i(M-1)\nu}\right)\right)^n\right] \\ &= \mathbb{E}\left[\left(1-\frac{1}{M}\left(\sum_{i=1}^{M} e^{-L_i(M-1)\nu}\right)\right)^n \mathbf{1}\left(\bigcap_{i=1}^{M}\{L_i > x_\epsilon\}\right)\right] \\ &\quad + \mathbb{E}\left[\left(1-\frac{1}{M}\left(\sum_{i=1}^{M} e^{-L_i(M-1)\nu}\right)\right)^n \mathbf{1}\left(\bigcup_{i=1}^{M}\{L_i \le x_\epsilon\}\right)\right].\end{aligned}$$

Next, by using $1-x \le e^{-x}$ and the independence of L_i, we derive

$$\begin{aligned}\mathbb{P}[N>n] &\le \left(\mathbb{E}\left[e^{-\frac{n}{M}e^{-L(M-1)\nu}}\mathbf{1}(L>x_\epsilon)\right]\right)^M + \left(1-\frac{1}{M}e^{-x_\epsilon(M-1)\nu}\right)^n \\ &\le \left(\mathbb{E}\left[e^{-\frac{n}{M}e^{-L\mathbf{1}(L>x_\epsilon)(M-1)\nu}}\right]\right)^M + \eta^n, \end{aligned} \tag{4}$$

where $\eta \triangleq 1 - e^{-x_\epsilon(M-1)\nu}/M < 1$. Then, by assumption (1), for any $0<\epsilon<\mu$, we can choose x_ϵ such that $\mathbb{P}[L>x] \le e^{-(\mu-\epsilon)x}$ for all $x \ge x_\epsilon$, which, by defining an exponential random variable L_ϵ with $\mathbb{P}[L_\epsilon > x] = e^{-(\mu-\epsilon)x}, x \ge 0$, implies $L\mathbf{1}(L>x_\epsilon) \overset{d}{\le} L_\epsilon$, where " $\overset{d}{\le}$ " denotes inequality in distribution. Therefore, (4) implies

$$\mathbb{P}[N>n] \le \left(\mathbb{E}\left[e^{-\frac{n}{M}e^{-L_\epsilon(M-1)\nu}}\right]\right)^M + \eta^n. \tag{5}$$

Now, for any $0 < x < 1$,

$$\mathbb{P}\left[e^{-(\mu-\epsilon)L_\epsilon} < x\right] = \mathbb{P}[(\mu-\epsilon)L_\epsilon > -\log x] = x,$$

implying that $e^{-(\mu-\epsilon)L_\epsilon} \stackrel{d}{=} U$, where " $\stackrel{d}{=}$ " denotes equality in distribution and U is a uniform random variable between 0 and 1. Thus,

$$\mathbb{P}[N > n] \leq \left(\mathbb{E}\left[e^{-\frac{n}{M}U^{(M-1)\nu/(\mu-\epsilon)}}\right]\right)^M + \eta^n.$$

By using the identity $\mathbb{E}[e^{-\theta U^{1/\alpha}}] = \Gamma(\alpha+1)/\theta^\alpha$, one can easily obtain

$$\overline{\lim_{n\to\infty}} \frac{\log \mathbb{P}[N > n]}{\log n} \leq -\frac{M(\mu-\epsilon)}{(M-1)\nu}, \tag{6}$$

which, by passing $\epsilon \to 0$, completes the proof of the upper bound.

For the lower bound, define $L_o \triangleq \min\{L_1, L_2, \cdots, L_M\}$, and observe that

$$\begin{aligned}\mathbb{P}[N > n] &= \mathbb{E}\left[\left(1 - \frac{1}{M}\left(\sum_{i=1}^{M} e^{-L_i(M-1)\nu}\right)\right)^n\right] \\ &\geq \mathbb{E}\left[\left(1 - e^{-L_o(M-1)\nu}\right)^n\right].\end{aligned} \tag{7}$$

The complementary cumulative distribution function $\bar{F}_o(x) \triangleq \mathbb{P}[L_o \geq x]$ satisfies

$$\lim_{x\to\infty} \frac{\log \bar{F}_o(x)}{x} = -M\mu,$$

implying that, for any $\epsilon > 0$, there exists x_ϵ such that $\mathbb{P}[L_o > x] \geq e^{-(M\mu+\epsilon)x}$ for all $x \geq x_\epsilon$. Next, if we define random variable L_o^ϵ such that $\mathbb{P}[L_o^\epsilon > x] = e^{-(M\mu+\epsilon)x}, x \geq 0$, then,

$$L_o \stackrel{d}{\geq} L_o^\epsilon \mathbf{1}(L_o^\epsilon > x_\epsilon),$$

which, by (7), implies

$$\begin{aligned}\mathbb{P}[N > n] &\geq \mathbb{E}\left[\left(1 - e^{-L_o^\epsilon \mathbf{1}(L_o^\epsilon > x_\epsilon)(M-1)\nu}\right)^n\right] \\ &\geq \mathbb{E}\left[\left(1 - e^{-L_o^\epsilon(M-1)\nu}\right)^n \mathbf{1}(L_o^\epsilon > x_\epsilon)\right].\end{aligned}$$

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

$$\begin{aligned}\mathbb{P}[N > n] &\geq \mathbb{E}\left[e^{-(1-\epsilon)ne^{-L_o^\epsilon(M-1)\nu}} \mathbf{1}(L_o^\epsilon > x_\epsilon)\right] \\ &\geq \mathbb{E}\left[e^{-(1-\epsilon)ne^{-L_o^\epsilon(M-1)\nu}}\right] - \mathbb{E}\left[e^{-(1-\epsilon)ne^{-L_o^\epsilon(M-1)\nu}} \mathbf{1}(L_o^\epsilon \leq x_\epsilon)\right] \\ &\geq \mathbb{E}\left[e^{-(1-\epsilon)ne^{-L_o^\epsilon(M-1)\nu}}\right] - \zeta^n,\end{aligned} \tag{8}$$

where $\zeta = e^{-(1-\epsilon)e^{-x_\epsilon(M-1)\nu}} < 1$. Similarly as in the proof of the upper bound, it is easy to check that $e^{-(M\mu+\epsilon)L_o^\epsilon} \overset{d}{=} U$, and therefore,

$$\mathbb{P}[N > n] \geq \mathbb{E}\left[e^{-(1-\epsilon)nU^{(M-1)\nu/(M\mu+\epsilon)}}\right] - \zeta^n,$$

which, by recalling the identity $\mathbb{E}[e^{-\theta U^{1/\alpha}}] = \Gamma(\alpha+1)/\theta^\alpha$, yields

$$\varliminf_{n\to\infty} \frac{\log \mathbb{P}[N > n]}{\log n} \geq -\frac{M\mu+\epsilon}{(M-1)\nu}. \tag{9}$$

Finally, passing $\epsilon \to 0$ in (9) completes the proof of the lower bound. Combining the lower and upper bound, we finish the proof of (2) for the case $U(0) = M$.

Next, define $N_s \triangleq \min\{n \geq 0 : U(T_n) = M\}$, $N_l \triangleq \min\{N, N_s\}$ and $N_e \triangleq N - N_l$. It can be shown that

$$\mathbb{P}[N_l > n] \leq \mathbb{P}[N_s > n] = o(e^{-\theta n}) \tag{10}$$

for some $\theta > 0$; due to space limits, the details of this proof will be presented in the full version of this paper. Assuming that the preceding bound holds, by the memoryless property of exponential distributions, we obtain

$$\begin{aligned}\mathbb{P}[N_e > n] &= \mathbb{P}[N - N_s > n, N > N_s]\\ &= \mathbb{P}[N > N_s]\mathbb{P}[N - N_s > n \mid N > N_s]\\ &= \mathbb{P}[N > N_s]\mathbb{P}[N > n \mid N_s = 0].\end{aligned}$$

Noting that $\mathbb{P}[N > n \mid N_s = 0]$ is the case of $U(0) = M$ that has already been proved, we conclude

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[N_e > n]}{\log n} = -\frac{M\mu}{(M-1)\nu},$$

which, combined with (10) and the union bound, yields

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[N > n]}{\log n} = \lim_{n\to\infty} \frac{\log \mathbb{P}[N_l + N_e > n]}{\log n} = -\frac{M\mu}{(M-1)\nu}. \tag{11}$$

The proof of (3) is presented in Section 4. □

3 Power Laws in Slotted ALOHA with Random Number of Users

It is clear from the preceding section that the power law delays arise due to the combination of collisions and packet variability. Hence, it is reasonable to expect an improved performance when this variability is reduced. Indeed, it is easy to see that the delays are geometrically bounded in a slotted ALOHA with constant size packets and a finite number of users. However, in this section we will show that, when the number of users sharing the channel has asymptotically an exponential distribution, the slotted ALOHA exhibits power law delays as well. Situations with random number of users are essentially predominant in practice, e.g., in sensor networks, the number of active sensors in a neighborhood is a random variable since sensors may switch between sleep/active modes, as shown in Figure 2; similarly in ad hoc wireless networks the variability of users may arise due to mobility, new users joining the network, etc.

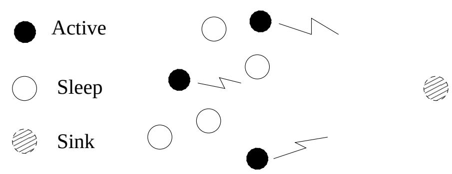

Fig. 2. Random number of active neighbors in a sensor network.

More formally, consider a slotted ALOHA model (e.g., see Section 4.2.2 of [3]) with packets/slots of unit size and a random number of users $M \geq 1$ that are fixed over time. Similarly as in Section 2, each user holds at most one packet at a time and after a successful transmission a new packet is generated according to an independent Bernoulli process with success probability $0 < \lambda \leq 1$. In case of a collision, each colliding user backs off according to an independent geometric random variable with parameter $e^{-\nu}, \nu > 0$. Denote the number of slots where transmissions are attempted but failed and the total time between two successful packet transmissions as N and T, respectively.

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

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

then, we have

$$\lim_{n \to \infty} \frac{\log \mathbb{P}[N > n]}{\log n} = \lim_{t \to \infty} \frac{\log \mathbb{P}[T > t]}{\log t} = -\frac{\alpha}{\nu}. \tag{12}$$

Remark 3. Similarly as in Theorem 1, this result shows that the distributions of N and T are essentially power laws, i.e., $\mathbb{P}[T > t] \approx t^{-\alpha/\nu}$ and, clearly, if $\alpha < \nu$, then $\mathbb{E}N = \mathbb{E}T = \infty$.

Proof. First consider a situation where all the users are backlogged, i.e., have a packet to send. In this case the total number of collisions between two successful transmissions is geometrically distributed given M,

$$\mathbb{P}[N > n \mid M] = \left(1 - \frac{Me^{-(M-1)\nu}(1 - e^{-\nu})}{1 - e^{-M\nu}}\right)^n, \ n \in \mathbb{N},$$

since, given M, $1-e^{-M\nu}$ is the conditional probability that there is an attempt to transmit a packet, and $1-e^{-M\nu}-Me^{-(M-1)\nu}(1-e^{-\nu})$ is the conditional probability that there is a collision. Therefore,

$$\mathbb{P}[N>n]=\mathbb{E}\left[\left(1-\frac{Me^{-(M-1)\nu}(1-e^{-\nu})}{1-e^{-M\nu}}\right)^n\right]. \tag{13}$$

On the other hand, we have

$$\mathbb{P}[T>t]=\mathbb{E}\left[\left(1-Me^{-(M-1)\nu}(1-e^{-\nu})\right)^t\right],\ t\in\mathbb{N}. \tag{14}$$

Now, following the same arguments as in the proof of Theorem 1, we can prove (12). Similarly, one can show that the same asymptotic results hold if the initial number of backlogged users is less than M. Due to space limitations, a complete proof of this theorem will be presented in the extended version of this paper. □

Actually, using the technique developed in [6] with some modifications, we can compute the exact asymptotics under a bit more restrictive conditions. Again, due to space limitations, the proof of the following theorem is deferred to the full version of the paper.

Theorem 3. *If $\lambda=\nu$ and $\bar{F}(x)\triangleq\mathbb{P}[M>x]$ satisfies $H\left(-\log\bar{F}(x)\right)\bar{F}(x)^{1/\beta}\sim xe^{-\nu x}$ with $H(x)$ being continuous and regularly varying, then, as $t\to\infty$,*

$$\mathbb{P}[T>t]\sim\frac{\Gamma(\beta+1)(e^\nu-1)^\beta}{t^\beta H(\beta\log t)^\beta}.$$

4 Simulation Examples

In this section, we illustrate our theoretical results with simulation experiments. In particular, we emphasize the characteristics of the studied ALOHA protocol that may not be immediately apparent from our theorems. For example, in practice, the distributions of packets and number of random users might have a bounded supports. We show that this situation may result in truncated power law distributions for T. 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 M and, thus, may result in very long, although, exponentially bounded delays. We will study the case where M has a bounded support in our second experiment.

Example 1 (Finite population model). For the finite population model described in Subsection 2.1, we study the distribution of time T between two consecutive successful transmissions. In this regard, we conduct four experiments for $M=2,4,10,20$ users, respectively. The packets are assumed i.i.d. exponential with mean 1 and the arrival intervals and backoffs follow exponential distribution with mean 2/3. Simulation result with 10^5 samples are shown in Figure 3, which

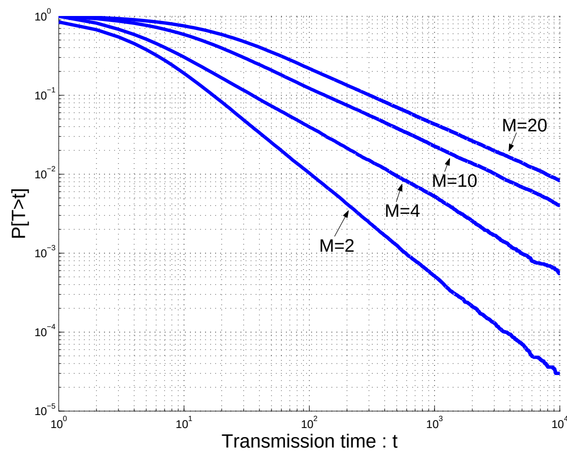

Fig. 3. Interval distribution between successfully transmitted packets for finite population ALOHA with variable size packets.

indicates a power law transmission delay. We can see from the figure that, as M gets large ($M = 10, 20$), the slopes of the distributions that represent the power law exponents on the log / log plot are essentially the same, as predicted by our Theorem 1.

Example 2 (Random number of users). As stated in Section 3, the situation when the number of users M is random may cause heavy-tailed transmission delays even for slotted ALOHA. However, in many practical applications the number of active users M may be bounded, i.e., the distribution $\mathbb{P}[M > x]$ has a bounded support. Thus, from equation (14) it is easy to see that the distribution of T is exponentially bounded. However, this exponential behavior may happen for very small probabilities, while the delays of interest can fall inside the region of the distribution (main body) that behaves as the power law. This example is aimed to illustrate this important phenomenon. Assume that initially $M \geq 1$ users have unit size packets ready to send and M follows geometric distribution with mean 3. The backoff times of colliding users are independent and geometrically distributed with mean 2. We take the number of users to have finite support $[1, K]$ and show how this results in a truncated power law distribution for T in the main body, even though the tails are exponentially bounded. This example is parameterized by K where K ranges from 6 to 14 and for each K we set the number of users to be equal to $M_K = \min(M, K)$. We plot the distribution of $\mathbb{P}[T > t]$, parameterized by K, in Figure 4. From the figure we can see that, when we increase the support of the distributions from $K = 6$ to $K = 14$, the main (power law) body of the distribution of T increases from less than 5 to almost 700. This effect is what we call the stretched support of the main body of $\mathbb{P}[T > t]$ in relation to the support K of M. In fact, it can be rigorously shown that the support of the main body of $\mathbb{P}[T > t]$ grows

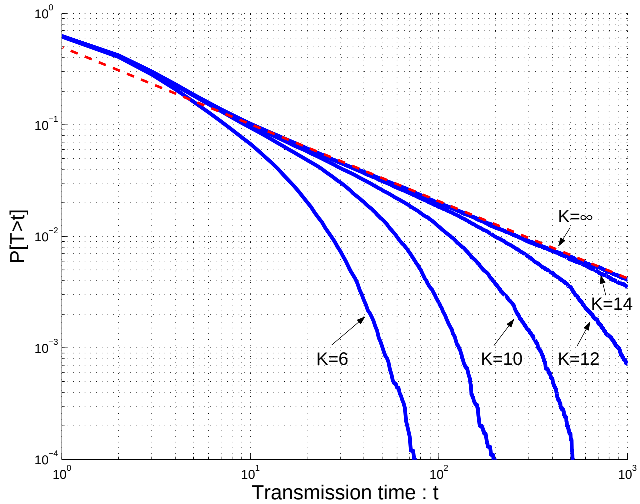

Fig. 4. Illustration of the stretched support of the power law main body when the number of users is $\min(M, K)$, where M is geometrically distributed.

exponentially fast. Furthermore, it is important to note that, if $K = 14$ and the probabilities of interest for $\mathbb{P}[T > t]$ are bigger than 1/500, then the result of this experiment is basically the same as for $K = \infty$; see Figure 4.

Appendix

Proof (of equation (3)). Similarly as in proving (2), we first assume that at time $t = 0$ there is a collision and that $U(0) = M$. Next, let $\{Y_i\}_{i \geq 1}$ be an i.i.d. sequence of exponential random variables with parameter $M\mu$ and define $L^* = \max_{1 \leq i \leq M}\{L_i\}$ where L_i is the packet size. Then,

$$\sum_{i=1}^{N} Y_i \overset{d}{\leq} T \overset{d}{\leq} NL^* + \sum_{i=1}^{N} Y_i. \tag{15}$$

Now, we establish the upper bound. For $H > 0$,

$$\begin{aligned}\mathbb{P}[T > t] &\leq \mathbb{P}\left[N > \frac{t}{H \log t}\right] + \mathbb{P}\left[\sum_{1 \leq i \leq t/H \log t} Y_i > \frac{t}{2}\right] + \mathbb{P}\left[\frac{tL^*}{H \log t} > \frac{t}{2}\right] \\ &\triangleq I_1 + I_2 + I_3. \end{aligned} \tag{16}$$

Since L^* is exponentially bounded, we can choose H large enough, such that for any fixed $\alpha > 0$,

$$I_3 = o\left(\frac{1}{t^\alpha}\right). \tag{17}$$

For the second term, by applying union bound, we obtain

$$I_2 \leq \left(\frac{t}{H\log t}+1\right)\mathbb{P}\left[Y_i > \frac{H\log t}{2}\right] = o\left(\frac{1}{t^{\alpha}}\right), \tag{18}$$

for any fixed $\alpha > 0$ and H large enough. Next, by (2), the asymptotics of I_1 is equal to

$$\lim_{t\to\infty}\frac{\log\mathbb{P}\left[N > \frac{t}{H\log t}\right]}{\log t} = -\frac{M\mu}{(M-1)\nu},$$

which, combined with (16), (17) and (18), completes the proof of the upper bound. Next, we prove the the lower bound. By the left inequality of (15), for $1 > \delta > 0$,

$$\mathbb{P}[T > t] \geq \mathbb{P}\left[N \geq \frac{t(1+\delta)}{\mathbb{E}[Y_1]}+1\right] - \mathbb{P}\left[\sum_{i=1}^{N} Y_i \leq t, N \geq \frac{t(1+\delta)}{\mathbb{E}[Y_1]}+1\right]. \tag{19}$$

Now, by using the standard large deviation result (Chernoff bound), it immediately follows that, for some $\eta > 0$, the second term on the right-hand side of (19) is bounded by

$$\mathbb{P}\left[\sum_{i\leq t(1+\delta)/\mathbb{E}[Y_1]+1} Y_i \leq t\right] = \mathbb{P}\left[\sum_{i\leq t(1+\delta)/\mathbb{E}[Y_1]+1}(\mathbb{E}[Y_1]-Y_i) \geq \delta t\right] = o\left(e^{-\eta t}\right).$$

Again, by (2), the first term on the right-hand side of (19) gives the right asymptotics, which proves the lower bound. The proof of the case $U(0) < M$ is more involved and, therefore, due to space limitation, we defer it to the full version of this paper. □

References

1. N. Abramson. The Aloha system - another alternative for computer communications. *Proc. of the Fall Joint Computer Conference*, pages 281–285, 1970.
2. S. Bellini and F. Borgonovo. On the throughput of an ALOHA channel with variable length packets. *IEEE Transactions on Communications*, 28:1932–1935, Nov. 1980.
3. D. P. Bertsekas and R. Gallager. *Data Networks.* Prentice Hall, 2nd edition, 1992.
4. M. J. Ferguson. An approximate analysis of delay for fixed and variable length packets in an unslotted ALOHA channel. *IEEE Transactions on Communications*, 25:644–654, July 1977.
5. 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*, 33(2):15–17, September 2005.
6. P. R. Jelenković and J. Tan. Can retransmissions of superexponential documents cause subexponential delays? In *Proceedings of IEEE INFOCOM'07*, Anchorage, Alaska, USA, May 2007. URL: http://www.comet.columbia.edu/~predrag/pub_list.html.

7. V. Naware, G. Mergen, and L. Tong. Stability and delay of finite user slotted ALOHA with multipacket reception. 51(7):2636–2656, July 2005.
8. R. Rom and M. Sidi. *Multiple access protocols: performance and analysis*. Springer-Verlag New York, Inc., New York, NY, USA, 1990.
9. M. Schwartz. *Telecommunication networks: protocols, modeling and analysis*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1986.
10. 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.