HEAVY TAILS IN QUEUEING SYSTEMS: IMPACT OF PARALLELISM ON TAIL PERFORMANCE

BO JIANG,* ** *University of Massachusetts*

JIAN TAN,*** *The Ohio State University and IBM T. J. Watson Research*

WEI WEI,* **** *University of Massachusetts*

NESS SHROFF,***** *The Ohio State University*

DON TOWSLEY,* ****** *University of Massachusetts*

Abstract

In this paper we quantify the efficiency of parallelism in systems that are prone to failures and exhibit power law processing delays. We characterize the performance of two prototype schemes of parallelism, *redundant* and *split*, in terms of both the power law exponent and exact asymptotics of the delay distribution tail. We also develop the optimal splitting scheme which ensures that split always outperforms redundant.

Keywords: Multipath; power law; parallelism

2010 Mathematics Subject Classification: Primary 68M20
Secondary 60G99

1. Introduction and model description

Parallelism is a common approach to improve reliability and efficiency in practice. Of all the diverse forms of parallelism, two prototype schemes stand out: *redundant* and *split*. In the *redundant* scheme a task is processed in its entirety by each agent, and is considered as completed when any one of the agents finishes. In the *split* scheme a task is split into multiple subtasks, each processed independently by a different agent, and the original task is completed when all subtasks are. In both cases, we expect better efficiency from using parallelism either because the processing time is the minimum of all the agents or because a smaller task needs to be completed by each agent.

In this paper we quantify the efficiency of parallelism in mitigating power law tails, which have been shown to be present when a job needs to be restarted after a failure occurs [3], [9], [10], [11], [14]. For the sake of definiteness, let us consider the notion of parallelism in the context of communication networks, where a data unit can be transmitted using multiple paths. A data unit can be a file or a packet (which are henceforth used interchangeably), and the
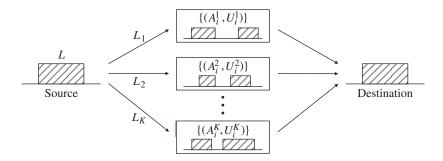

FIGURE 1: Multipath transmission over K channels with failures.

transmission needs to restart after a failure (i.e. there is no check point in the transmission). In Figure 1 we show a sketch of the multipath model considered in this paper, which is a generalization of the single-path model introduced in [10]. There are K independent paths between the source and the destination. The channel dynamics of path j, $1 \leq j \leq K$, are modeled as an on–off process $\{(A_i^j, U_i^j)\}_{i \geq 1}$ that alternates between available periods A_i^j and unavailable periods U_i^j. We assume that $\{A_i^j\}_{i \geq 1}$ are independent and identically distributed (i.i.d.) with common distribution A^j, and $\{U_i^j\}_{i \geq 1}$ are i.i.d. with common distribution U^j. Moreover, the sequences $\{U_i^j\}_{i \geq 1}$ and $\{A_i^j\}_{i \geq 1}$, $1 \leq j \leq K$, are mutually independent.

Let L be the random variable denoting the length of a packet, which is assumed to be independent of the channel dynamics, i.e. $\{(A_i^j, U_i^j)\}_{i \geq 1}$. A fragment of length $L_j = \gamma_j L$, $0 \leq \gamma_j \leq 1$, of the packet is sent over path j. Packet transmissions can start only at the beginnings of available periods. A transmission over path j that starts at the beginning of A_i^j is considered successful if $A_i^j \geq L_j$; otherwise, the transmission aborts and waits for the beginning of the next available period A_{i+1}^j to restart.

We study two multipath transmission schemes, namely, *redundant transmission* and *split transmission*, corresponding to the two aforementioned prototypes of parallelism. Under redundant transmission, the same packet is transmitted in its entirety over all K paths, so $\gamma_j = 1$ for all j, and the transmission is successful as soon as one of the K replicas arrives at the destination. Under split transmission, the packet is split into K nonoverlapping fragments, each sent over a different path, so $\sum_{1 \leq j \leq K} \gamma_j = 1$, and the transmission is complete only when the last fragment arrives at the destination. The quantity of interest is the overall transmission delay, of which the precise definition is given below.

Definition 1.1. The number of (re)transmissions of a packet of length L_j over path j, $1 \leq j \leq K$, is defined as

$$N_j := \min\{i : A_i^j \geq L_j\},$$

and the corresponding transmission delay over this path is defined as

$$T_j := \sum_{i=1}^{N_j - 1} (A_i^j + U_i^j) + L_j.$$

- Redundant transmission ($L_j = L$): the transmission is complete when the packet is successfully transmitted over any one of the K paths. Therefore, the overall transmission delay T^R is

$$T^R := \min_{1 \leq j \leq K} T_j.$$

- Split transmission ($\sum_{j=1}^{K} L_j = L$): the transmission is complete when all K fragments of the packet are successfully transmitted. Therefore, the overall transmission delay T^S is

$$T^S := \max_{1 \le j \le K} T_j,$$

and the total number of retransmissions over the K paths is

$$N^S := \sum_{j=1}^{K} N_j.$$

Our main contributions in this paper can be summarized as follows.

- We characterize the asymptotic behaviors of $\mathbb{P}[T^R > x]$ and $\mathbb{P}[T^S > x]$, in terms of both the power law exponent (Theorems 3.2 and 4.2) and exact asymptotics (Theorems 3.3 and 4.3). Compared to the single-path transmission on the best path, redundant transmission does not change the power law tail exponent of the delay distribution (Theorem 3.2), but only decreases the distribution tail by a constant factor (Theorem 3.3). On the other hand, depending on the packet size distribution and the manner of splitting, split transmission could either increase or decrease the power law tail exponent (Theorem 4.2).

- We develop the optimal split transmission scheme that minimizes the power law tail exponent of the transmission delay, which is guaranteed to be no larger than that of redundant transmission and the best single-path transmission (Theorem 4.4). The optimal split transmission scheme is effective in mitigating power law delays if the absolute value of the logarithm of the packet size probability tail is regularly varying with positive index, and becomes ineffective if the above quantity is slowly varying.

Multipath transmissions have also been studied in [1] using extreme value theory, with the number of paths going to infinity. In the present work we focus on the context of multipath transmissions in communication networks with a fixed (typically small) number of paths, where the multipath transmission has long been used to improve reliability and efficiency (see, e.g. [5], [6], and [12]). Here we want to emphasize that the packet size distribution has been assumed to have an infinite support in this study, which contradicts the reality that all packet networks (from the Internet to wireless LANs) impose maximum packet sizes at different layers of the protocol stack. It can be proved that eventually the transmission delay distribution will be light tailed under this condition. However, as has been shown in [15], this light-tailed behavior occurs with a power law main body of the delay distribution, and this power law behavior may have dominating effects on the system performance since it spans over a time interval that increases very fast with respect to the length of the longest packet. Thus, our assumption on the infinite support of the packet size distribution allows us to study the main body of the transmission delay distribution. While, similar to [15], we can extend our results to the case with packets having finite support, we feel that this would distract from the main insights gained from the paper.

Note also that while we have chosen to cast the mathematical model in the context of data transmission for wireless networks, especially for low-power sensor networks where simple operations are preferred to recover failed data (for the performance with complicated coding schemes, see [16]), the model is applicable to many other scenarios that involve parallelism and job failures, such as computing jobs in grid computing, file downloading in peer-to-peer networks, parallel experiment planning, and parallel scheduling.

The rest of the paper is organized as follows. In Section 2 we summarize the known results on single-path transmission. Redundant transmission is investigated in Section 3, and split transmission in Section 4.

2. Summary of known results on single-path transmission

In this section we establish some notation that will be used throughout the paper, and also summarize the results on single-path transmission that will be used later.

Throughout the paper, we will use the following notation to denote the complementary cumulative distribution functions of A^j, $1 \le j \le K$, and L:

$$\bar{G}_j(x) := \mathbb{P}[A^j > x]$$

and

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

with $\bar{F}(x)$ being continuous eventually. The K paths are said to be *homogeneous* if $\{A^j, U^j\}_{1 \le j \le K}$ are identically distributed as $\{A, U\}$, in which case we use $\bar{G}(x) := \mathbb{P}[A > x]$. In general, $\{A^j, U^j\}_{1 \le j \le K}$ are not identically distributed, and the K paths are said to be *heterogeneous*.

We will use the limit

$$\alpha_j := \lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)},$$

when it exists, as a coarse quality measure of channel j relative to the packet size distribution, with a larger value corresponding to a better channel.

We will also assume some moment conditions on $\{U^j\}_{j=1}^K$, $\{A^j\}_{j=1}^K$, and L. Specifically, we will say that the *moment conditions hold with parameter α* if there exists some $\theta > 0$ such that

(C1) $\max_{1 \le j \le K} \mathbb{E}[(U^j)^{(\alpha \vee 1)+\theta}] < \infty$;

(C2) $\max_{1 \le j \le K} \mathbb{E}[(A^j)^{1+\theta}] < \infty$;

(C3) $\mathbb{E}[L^{\alpha+\theta}] < \infty$.

Before we proceed, recall the following definition of the *regularly varying function* [4].

Definition 2.1. A positive measurable function f is called regularly varying (at infinity) with index ρ if

$$\lim_{x \to \infty} \frac{f(\lambda x)}{f(x)} = \lambda^\rho \tag{2.1}$$

for all $\lambda > 0$. It is called slowly varying if $\rho = 0$.

Also, recall the standard definition of an inverse function $f^{\leftarrow}(x) := \inf\{y \colon f(y) > x\}$ for a nondecreasing function $f(x)$; note that the notation $f(x)^{-1}$ represents $1/f(x)$. We will use '$\vee$' to denote max, i.e. $x \vee y := \max\{x, y\}$. For any two real functions $f(x)$ and $g(x)$, the following standard notation will also be used:

- $f(x) \sim g(x)$ if and only if $\lim_{x \to \infty} f(x)/g(x) = 1$;

- $f(x) = o(g(x))$ if and only if $\lim_{x \to \infty} f(x)/g(x) = 0$;

- $f(x) = O(g(x))$ if and only if $\overline{\lim}_{x \to \infty} f(x)/g(x) < \infty$.

2.1. Single-path transmission

For the case $K = 1$, i.e. there is only a single path in the system, the total number of transmissions $N = N_1$ and transmission delay $T = T^R = T^S$ have been studied in [3], [10], and [11]. Below we quote Propositions 2.1 and 2.2 from [10] and [11], which show that both N and T can follow power law distributions regardless of how heavy or light the tails of A and L might be.

Proposition 2.1. *Suppose that*

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \alpha > 0.$$

Then

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

If, in addition, the moment conditions hold with parameter α then

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

Proposition 2.2. *Suppose that*

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

where $\Phi(\cdot)$ is regularly varying with index $\alpha > 0$. Then, as $n \to \infty$,

$$\mathbb{P}[N > n] \sim \frac{\Gamma(\alpha + 1)}{\Phi(n)}.$$

If, in addition, the moment conditions hold with parameter α then, as $t \to \infty$,

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

Note that $\bar{F}(x)^{-1} \sim \Phi(\bar{G}(x)^{-1})$ implies that $\lim_{x \to \infty} \log \bar{F}(x)/\log \bar{G}(x) = \alpha$ by Theorem 1.4.1 and Proposition 1.3.6 of [4]. Thus, Proposition 2.2 provides more refined results than Proposition 2.1 under more restrictive conditions. As mentioned in the introduction, the results in the preceding two propositions as well as those in the rest of the paper can be readily extended to the case where packet sizes are bounded, using similar techniques as in [15].

3. Redundant transmission

In this section we study the redundant transmission scheme. We investigate whether redundant transmission over K paths can mitigate the power-law-distributed transmission delay suffered by single-path transmissions. We begin with the special case of homogeneous paths, followed by the general case of heterogeneous paths.

3.1. Homogeneous paths

In this section we present the results for homogeneous paths. We first consider the case where all packets are of the same size, and then the more realistic case where packet sizes are variable.

Proposition 3.1. *Suppose that all packets are of the same size l, and that $U = 0$. Then*

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^R > t]}{t} = -K\gamma,$$

where $\gamma > 0$ is the solution to the equation $\int_0^l e^{\gamma x} \, d\mathbb{P}[A \le x] = 1$.

This result can be derived using Corollary 3.2 of [3]. It shows that redundant transmission greatly improves the performance when all packets are equally sized. As K increases, we obtain order gains in the decay rate of the delay distribution tail.

In reality, however, packets are not equally sized due to many other considerations, e.g. reducing communication costs and extra overhead induced from encapsulation. We now present a theorem for the case where the packet size is a random variable.

Theorem 3.1. *Suppose that*

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \alpha > 0,$$

and that the moment conditions hold with parameter α. Then

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

Comparing the above theorem with Proposition 2.1, we observe that the power law tail exponent of the delay distribution under redundant transmission is the same as that under single-path transmission. This is because the packets sent over these paths are replicas of each other and, hence, $T_1, \ldots, T_K$ are not independent. This theorem is a direct consequence of Theorem 3.2 below, which investigates a more general scenario.

3.2. Heterogeneous paths

For heterogeneous paths, we have the following result for redundant transmission.

Theorem 3.2. *Suppose that*

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \alpha_j > 0, \qquad j = 1, 2, \ldots, K.$$

Let $\alpha^* := \max_{1 \le j \le K} \alpha_j > 0$ and $\Omega_K^* = \{j \in \{1, 2, \ldots, K\} : \alpha_j = \alpha^*\}$. If the moment conditions hold with parameter α^* and with (C1) replaced by

(C1′) $\min_{j \in \Omega_K^*} \mathbb{E}[(U^j)^{(\alpha^* \vee 1) + \theta}] < \infty$,

then

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^R > t]}{\log t} = -\alpha^*. \tag{3.1}$$

Theorem 3.2 shows that the tail behavior of the delay distribution under redundant transmission is determined by the best paths (i.e. the paths with the largest α_j).

Proof of Theorem 3.2. We first establish an upper bound. Suppose that path k achieves the minimum in (C1′). Note that $T_k \ge T^R = \min_{1 \le j \le K} T_j$. By Proposition 2.1,

$$\varlimsup_{t \to \infty} \frac{\log \mathbb{P}[T^R > t]}{\log t} \le \lim_{t \to \infty} \frac{\log \mathbb{P}[T_k > t]}{\log t} = -\alpha_k = -\alpha^*. \tag{3.2}$$

Next, we establish a lower bound by constructing a new system with longer available periods than all of the K paths. The new system has an on–off channel characterized by alternating i.i.d. sequences $\{\bar{A}_i\}$ and $\{\bar{U}_i\}$, where

$$\bar{A}_i = \max_{1 \leq j \leq K} A_i^j \quad \text{and} \quad \bar{U}_i = 0.$$

Denote by $\underline{N}$ the number of transmissions for a packet of length L over this newly constructed channel. Note that $\underline{N} \leq \min_{1 \leq j \leq K} N_j$.

Note that

$$\{A_i^j > x\} \subset \{\bar{A}_i > x\} = \bigcup_{j=1}^{K} \{A_i^j > x\}.$$

The monotonicity of the probability measure and the union bound yield

$$\max_{1 \leq j \leq K} \bar{G}_j(x) = \max_{1 \leq j \leq K} \mathbb{P}[A_i^j > x] \leq \mathbb{P}[\bar{A}_i > x] \leq \sum_{j=1}^{K} \mathbb{P}[A_i^j > x] \leq K \max_{1 \leq j \leq K} \bar{G}_j(x).$$

Thus, for x large enough so that $K \max_{1 \leq j \leq K} \bar{G}_j(x) < 1$, we have

$$\max_{1 \leq j \leq K} \frac{\log \bar{F}(x)}{\log K + \log \bar{G}_j(x)} \leq \frac{\log \bar{F}(x)}{\log \mathbb{P}[\bar{A}_i > x]} \leq \max_{1 \leq j \leq K} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)}.$$

Letting $x \to \infty$, we obtain

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \mathbb{P}[\bar{A}_i > x]} = \alpha^*.$$

Since $\mathbb{E}[(\bar{A}_i)^{1+\theta}] \leq \sum_{j=1}^{K} \mathbb{E}[(A_i^j)^{1+\theta}] < \infty$, Proposition 2.1 yields

$$\lim_{n \to \infty} \frac{\log \mathbb{P}[\underline{N} > n]}{\log n} = -\alpha^*. \tag{3.3}$$

Now define $\underline{A}_i = \min_{1 \leq j \leq K} A_i^j$. Note that $T^{\mathrm{R}} \geq \sum_{i=1}^{N-1} \underline{A}_i$, so

$$\mathbb{P}[T^{\mathrm{R}} > t] \geq \mathbb{P}\left[\sum_{i=1}^{N-1} \underline{A}_i > t\right]$$

$$\geq \mathbb{P}\left[\sum_{i=1}^{\lfloor t \log t \rfloor} \underline{A}_i > t, \ \underline{N} > t \log t\right]$$

$$\geq \mathbb{P}[\underline{N} > t \log t] - \mathbb{P}\left[\sum_{i=1}^{\lfloor t \log t \rfloor} \underline{A}_i \leq t\right], \tag{3.4}$$

where the first two inequalities follow from the monotonicity of the probability measure and the fact that $\underline{N} - 1 \geq \lfloor t \log t \rfloor$ for $\underline{N} > t \log t$, and the last inequality follows from $\mathbb{P}[A \cap B] \geq \mathbb{P}[A] - \mathbb{P}[B^c]$.

Using the Markov inequality and the fact that the $\underline{A}_i$s are i.i.d.,

$$\mathbb{P}\left[\sum_{i=1}^{\lfloor t\log t\rfloor}\underline{A}_i \le t\right] = \mathbb{P}\left[\exp\left(-\sum_{i=1}^{\lfloor t\log t\rfloor}\underline{A}_i\right) \ge \mathrm{e}^{-t}\right]$$

$$\le \frac{\mathbb{E}\exp(-\sum_{i=1}^{\lfloor t\log t\rfloor}\underline{A}_i)}{\mathrm{e}^{-t}}$$

$$= \mathrm{e}^t(\mathbb{E}\mathrm{e}^{-\underline{A}_1})^{\lfloor t\log t\rfloor}.$$

Since $\underline{A}_1 \ge 0$ and $\mathbb{P}[\underline{A}_1 > 0] > 0$, we have $0 < \mathbb{E}\mathrm{e}^{-\underline{A}_1} < 1$, so $\mathbb{P}[\sum_{i=1}^{\lfloor t\log t\rfloor}\underline{A}_i \le t]$ drops off exponentially in $t\log t$. On the other hand, (3.3) shows that $\mathbb{P}[\underline{N} > t\log t]$ drops off algebraically in $t\log t$, so (3.4) yields

$$\mathbb{P}[T^{\mathrm{R}} > t] \ge (1 + o(1))\mathbb{P}[\underline{N} > t\log t].$$

Noting that $\log(t\log t) \sim \log t$ and invoking (3.3) again, we obtain

$$\lim_{t\to\infty}\frac{\log\mathbb{P}[T^{\mathrm{R}} > t]}{\log t} \ge \lim_{t\to\infty}\frac{\log\mathbb{P}[\underline{N} > t\log t]}{\log(t\log t)} = -\alpha^*,$$

which, together with (3.2), establishes (3.1). This completes the proof of Theorem 3.2.

Theorem 3.2 characterizes the performance in terms of the logarithmic asymptotics. Basically, it only contains information about the power law tail exponent, but yields no information beyond. As a consequence, this result cannot distinguish between redundant transmission and single-path transmission over the best path(s). In order to investigate the performance gain for redundant transmission, we need a more refined asymptotic result. For a set of regularly varying functions $\Phi_j(\cdot)$, $1 \le j \le K$, we can compute the exact asymptotic tail of the distribution of T^{R}.

Theorem 3.3. *Suppose that $\bar{F}(x)^{-1} \sim \zeta_j\Phi_j(\bar{G}_j(x)^{-1})$, where $\zeta_j > 0$, and $\Phi_j(\cdot)$ is regularly varying with index $\alpha_j > 0$ such that $\Phi_i(x) \sim \Phi_j(x)$ if $\alpha_i = \alpha_j$. Let $\alpha^* = \max_{1\le j\le K}\alpha_j$ and $\Omega_K^* = \{j \in \{1, 2, \ldots, K\}: \alpha_j = \alpha^*\}$. If the moment conditions hold with parameter α^* and with (C1) replaced by*

(C1″) $\max_{j\in\Omega_K^*}\mathbb{E}[(U^j)^{(\alpha^*\vee 1)+\theta}] < \infty,$

then, as $t \to \infty$,

$$\mathbb{P}[T^{\mathrm{R}} > t] \sim \frac{\Gamma(\alpha^* + 1)}{(\sum_{j\in\Omega_K^*}(\mathbb{E}[A^j + U^j])^{-1}\zeta_j^{1/\alpha^*})^{\alpha^*}}\frac{1}{\Phi^*(t)},$$

where $\Phi^*(t) \sim \Phi_j(t)$ for $j \in \Omega_K^*$.

This result shows that, when there are multiple channels with the best quality measure α^*, redundant transmission improves the system performance by reducing the delay distribution tail by a constant factor, relative to the single-path transmission over any such path. Moreover, this constant factor does not depend on the nonbest paths. When the K channels are i.i.d., it is equal to K^α.

In order to prove the theorem, we need the following lemmas, which are stated for the general case where $L_j = \gamma_j L$ for some $\gamma_j > 0$, so that the results will be applicable later to the split transmission scheme. Recall that $\gamma_j = 1$ for redundant transmission and $\sum_{j=1}^K\gamma_j = 1$ for split transmission.

Lemma 3.1. *Suppose that* $\bar{F}(x)^{-1} \sim \zeta_j \Phi_j(\bar{G}_j(\gamma_j x)^{-1})$, *where* $\zeta_j > 0$, *and* $\Phi_j(\cdot)$ *is a regularly varying function with index* $\alpha_j > 0$ *such that* $\Phi_i(x) \sim \Phi_j(x)$ *if* $\alpha_i = \alpha_j$. *Then, for* $\psi_j > 0$, $j = 1, 2, \ldots, K$, *and a nonempty subset* $J \subset \{1, 2, \ldots, K\}$,

$$\mathbb{P}\left[\bigcap_{j \in J} \{N_j > \psi_j t\}\right] \sim \frac{\Gamma(\alpha_J^* + 1)}{(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*})^{\alpha_J^*}} \frac{1}{\Phi_J^*(t)}, \tag{3.5}$$

where $\alpha_J^* = \max_{j \in J} \alpha_j$, $J^* = \{j \in J : \alpha_j = \alpha_J^*\}$, *and* $\Phi_J^*(t) \sim \Phi_j(t)$ *for* $j \in J^*$.

Proof. See Appendix A.

Lemma 3.2. *Suppose that* $\mathbb{E}[(A^j)^{1+\theta}] < \infty$. *Then, for* $\psi_j > 1/\mathbb{E}[A^j]$, *there exists some* $\eta > 0$ *and* $C > 0$ *such that*

$$\mathbb{P}[T_j \leq t, \ N_j > \psi_j t] \leq C e^{-\eta t}.$$

If, in addition, $\mathbb{E}[(U^j)^{1+\theta}] < \infty$ *for some* $\theta > 0$ *then the claim is true for* $\psi_j > 1/\mathbb{E}[A^j + U^j]$.

Proof. Note that $N_j > \psi_j t$ implies that $N_j - 1 \geq \lfloor \psi_j t \rfloor$. Thus, for $N_j > \psi_j t$,

$$T_j = \sum_{i=1}^{N_j - 1} (A_i^j + U_i^j) + L_j \geq \sum_{i=1}^{\lfloor \psi_j t \rfloor} (A_i^j + U_i^j) \geq \sum_{i=1}^{\lfloor \psi_j t \rfloor} A_i^j,$$

from which it follows that

$$\mathbb{P}[T_j \leq t, \ N_j > \psi_j t] \leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi_j t \rfloor} (A_i^j + U_i^j) \leq t\right] \leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi_j t \rfloor} A_i^j \leq t\right].$$

By letting $X = A^j + U^j$, $X_i = A_i^j + U_i^j$, or $X = A^j$, $X_i = A_i^j$, we prove both cases at once. Given $y > 0$, the Markov inequality implies that

$$\mathbb{P}\left[\sum_{i=1}^{\lfloor \psi_j t \rfloor} X_i \leq t\right] = \mathbb{P}\left[\exp\left(-y \sum_{i=1}^{\lfloor \psi_j t \rfloor} X_i\right) \geq e^{-yt}\right] \leq e^{yt}(\mathbb{E}[e^{-yX}])^{\lfloor \psi_j t \rfloor}.$$

Choose $\delta > 0$ small enough so that $(1 - 2\delta)\psi_j \mathbb{E}X > 1$. Since $e^{-x} = 1 - x + o(x)$, there exists $x_0 > 0$ such that $e^{-x} \leq 1 - (1 - \delta)x$ for $0 \leq x \leq x_0$. Let $D = (1 - \delta)x_0^{-\theta} > 0$. Then, for $x \geq x_0$,

$$1 - (1 - \delta)x + Dx^{1+\theta} = 1 + (1 - \delta)x\left[\left(\frac{x}{x_0}\right)^{\theta} - 1\right] \geq 1 > e^{-x}.$$

Thus, $e^{-x} \leq 1 - (1 - \delta)x + Dx^{1+\theta}$ for all $x \geq 0$. Setting $x = yX$ and taking the expectation then yield, for small enough $y > 0$,

$$\mathbb{E}[e^{-yX}] \leq 1 - (1 - \delta)y\mathbb{E}X + Dy^{1+\theta}\mathbb{E}X^{1+\theta} \leq 1 - (1 - 2\delta)y\mathbb{E}X \leq e^{-(1-2\delta)y\mathbb{E}X}.$$

Therefore,

$$\mathbb{P}[T_j \leq t, \ N_j > \psi_j t] \leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi_j t \rfloor} X_i \leq t\right] \leq e^{yt}e^{-(1-2\delta)y\mathbb{E}X\lfloor \psi_j t \rfloor} = C e^{-\eta t},$$

where $\eta = y[(1 - 2\delta)\psi_j \mathbb{E}X - 1] > 0$ and $C = e^{(1-2\delta)y\mathbb{E}X}$.

Lemma 3.3. *If* $\mathbb{E}[(U^j)^{(\alpha \vee 1)+\theta}] < \infty$, $\mathbb{E}[(A^j)^{1+\theta}] < \infty$, $\mathbb{E}[L^{\alpha+\theta}] < \infty$ *for some* $\alpha > 0$ *and* $\theta > 0$, *and* $\psi_j < 1/\mathbb{E}[A^j + U^j]$, *then there exists* $\nu > \alpha$ *such that*

$$\mathbb{P}[T_j > t, \, N_j \le \psi_j t] = O(t^{-\nu}).$$

Proof. See Appendix B.

Now we prove Theorem 3.3.

Proof of Theorem 3.3. Let $\psi_j < 1/\mathbb{E}[A^j + U^j]$, $j \in \Omega_K^*$. Note that $\{T^{\mathrm{R}} > t\} \subset \bigcap_{j \in \Omega_K^*}\{T_j > t\}$. Thus,

$$\mathbb{P}[T^{\mathrm{R}} > t] \le \mathbb{P}\left[\bigcap_{j \in \Omega_K^*}\{T_j > t\}\right] \le \mathbb{P}\left[\bigcap_{j \in \Omega_K^*}\{N_j > \psi_j t\}\right] + \sum_{j \in \Omega_K^*}\mathbb{P}[T_j > t, \, N_j \le \psi_j t].$$

The last term is $o(1/\Phi^*(t))$ by Lemma 3.3 and Proposition 1.5.1 of [4]. Lemma 3.1 then yields

$$\lim_{t \to \infty} \Phi^*(t)\mathbb{P}[T^{\mathrm{R}} > t] \le \frac{\Gamma(\alpha^* + 1)}{(\sum_{j \in \Omega_K^*} \psi_j \zeta_j^{1/\alpha^*})^{\alpha^*}}. \tag{3.6}$$

Let $\widetilde{\psi}_j > 1/\mathbb{E}[A^j + U^j]$ for $j \in \Omega_K^*$ and $\widetilde{\psi}_j > 1/\mathbb{E}[A^j]$ for $j \notin \Omega_K^*$. Using union bounds, we obtain

$$\mathbb{P}[T^{\mathrm{R}} > t] \ge \mathbb{P}\left[\bigcap_{j=1}^K\{N_j > \widetilde{\psi}_j t\}\right] - \sum_{j=1}^K\mathbb{P}[T_j \le t, \, N_j > \widetilde{\psi}_j t].$$

The last term is $o(1/\Phi^*(t))$ by Lemma 3.2 and Proposition 1.5.1 of [4]. Lemma 3.1 then yields

$$\lim_{t \to \infty} \Phi^*(t)\mathbb{P}[T^{\mathrm{R}} > t] \ge \frac{\Gamma(\alpha^* + 1)}{(\sum_{j \in \Omega_K^*} \widetilde{\psi}_j \zeta_j^{1/\alpha^*})^{\alpha^*}}. \tag{3.7}$$

We complete the proof by combining (3.6) and (3.7) and letting $\psi_j, \widetilde{\psi}_j \to 1/\mathbb{E}[A^j + U^j]$ for $j \in \Omega_K^*$. □

4. Split transmission

In this section we study the split transmission scheme, where a packet is split into nonoverlapping fragments, each sent over a different path. Recall that a fraction γ_j of the packet L is sent over path j, where $\sum_{j=1}^K \gamma_j = 1$ and $0 \le \gamma_j \le 1$ for $1 \le j \le K$. We will assume that $\gamma_j > 0$ except in Theorem 4.4. We begin with the case of homogeneous paths, followed by the heterogeneous case. We also investigate which of the two schemes, split transmission or redundant transmission, results in a lighter tail for the transmission delay distribution. We develop the optimal splitting scheme that minimizes the tail exponent of the delay distribution, in which case split transmission outperforms redundant transmission.

4.1. Homogeneous paths

We have the following theorem for split transmission over homogeneous paths, where each packet is evenly split into K fragments. It is a special case of Theorem 4.2, so the proof is omitted.

Theorem 4.1. *Suppose that*

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \alpha > 0$$

and

$$\lim_{x \to \infty} \frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \beta. \tag{4.1}$$

If the moment conditions hold with parameter $\beta\alpha$ then

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^S > t]}{\log t} = -\beta\alpha.$$

Note that $\beta \geq 1$. By comparing Proposition 2.1 and Theorems 3.1 and 4.1, we observe that split transmission is no worse than redundant transmission for homogeneous paths, when packets are split evenly. Split transmission is not beneficial when $\beta = 1$, e.g. when $\log \bar{F}(x)^{-1}$ is a slowly varying function.

Theorem 4.1 shows that the effectiveness of split transmission is closely dependent on the packet size distribution, as characterized by (4.1). We illustrate this point further using several common distributions. For each distribution, we calculate α and β, and the power law tail exponent is $-\beta\alpha$.

Example 4.1. (*Weibull distribution.*) Suppose that both the packet size L and the available period A follow Weibull distributions, i.e.

$$\bar{F}(x) = \mathbb{P}[L > x] = e^{-(\lambda x)^b}, \qquad \bar{G}(x) = \mathbb{P}[A > x] = e^{-(\mu x)^b},$$

where $\lambda > 0$, $\mu > 0$, and $b > 0$. Then

$$\alpha = \lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \lim_{x \to \infty} \frac{\log e^{-(\lambda x)^b}}{\log e^{-(\mu x)^b}} = \left(\frac{\lambda}{\mu}\right)^b,$$

$$\beta = \lim_{x \to \infty} \frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \lim_{x \to \infty} \frac{\log (e^{-(\lambda Kx)^b})}{\log (e^{-(\lambda x)^b})} = K^b > 1.$$

Example 4.2. (*Pareto distribution.*) Suppose that both the packet size L and the available period A follow Pareto distributions, i.e.

$$\bar{F}(x) = \mathbb{P}[L > x] = \begin{cases} \left(\dfrac{b_0}{x}\right)^\lambda, & x \geq b_0, \\ 1, & x < b_0, \end{cases}$$

$$\bar{G}(x) = \mathbb{P}[A > x] = \begin{cases} \left(\dfrac{b_1}{x}\right)^\mu, & x \geq b_1, \\ 1, & x < b_1, \end{cases}$$

where $\lambda > 0$, $\mu > 0$, and $b_0, b_1 > 0$. Then

$$\alpha = \lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \lim_{x \to \infty} \frac{\lambda(\log b_0 - \log x)}{\mu(\log b_1 - \log x)} = \frac{\lambda}{\mu},$$

$$\beta = \lim_{x \to \infty} \frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \lim_{x \to \infty} \frac{\lambda(\log b_0 - \log K - \log x)}{\lambda(\log b_0 - \log x)} = 1.$$

Observe that $\beta = 1$ when L follows a Pareto distribution. In general, $\beta = 1$ if $\log \bar{F}(x)^{-1}$ is slowly varying. In that case, split transmission is not beneficial compared to single-path transmission and redundant transmission in terms of tail performance.

4.2. Heterogeneous paths

For heterogeneous paths, we have the following result for the transmission time.

Theorem 4.2. *Suppose that, for* $j = 1, 2, \ldots, K$,

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \alpha_j > 0, \tag{4.2}$$

$$\lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{F}(\gamma_j x)} = \beta_j. \tag{4.3}$$

Then

$$\lim_{n \to \infty} \frac{\log \mathbb{P}[N^S > n]}{\log n} = -\tau^\circ,$$

where $\tau^\circ = \min_{1 \le j \le K} \beta_j \alpha_j$. *If, in addition, the moment conditions hold with parameter* τ° *then*

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^S > t]}{\log t} = -\tau^\circ.$$

When paths are heterogeneous, the delay distribution tail is determined by the best path(s) under redundant transmission and by the worst path(s) under split transmission. On the other hand, split transmission only sends a fraction of the packet over each path. Comparing Theorems 3.2 and 4.2, we observe that, if $\min_{1 \le j \le K} \beta_j \alpha_j > \max_{1 \le j \le K} \alpha_j$, split transmission is more effective than redundant transmission in minimizing the power law tail exponent; otherwise, redundant transmission is more effective. We will show later that, by carefully choosing the way to split packets, the tail performance of split transmission is never worse than that of redundant transmission.

Proof of Theorem 4.2. We first prove the result for N^S. Since $\log \lfloor n/K \rfloor \sim \log n$ as $n \to \infty$, Proposition 2.1 then implies that

$$\lim_{n \to \infty} \frac{\log \mathbb{P}[N_j > n]}{\log n} = \lim_{n \to \infty} \frac{\log \mathbb{P}[N_j > n/K]}{\log n} = -\beta_j \alpha_j. \tag{4.4}$$

Since $N^S = \sum_{j=1}^K N_j$, we have

$$\max_{1 \le j \le K} \mathbb{P}[N_j > n] \le \mathbb{P}[N^S > n] \le \sum_{j=1}^K \mathbb{P}\left[N_j > \frac{n}{K}\right] \le K \max_{1 \le j \le K} \mathbb{P}\left[N_j > \frac{n}{K}\right],$$

which yields

$$-\tau^\circ = \max_{1 \le j \le K} \lim_{n \to \infty} \frac{\log \mathbb{P}[N_j > n]}{\log n}$$

$$\le \lim_{n \to \infty} \frac{\log \mathbb{P}[N^S > n]}{\log n}$$

$$\le \max_{1 \le j \le K} \lim_{n \to \infty} \frac{\log \mathbb{P}[N_j > n/K]}{\log n}$$

$$= -\tau^\circ, \quad \text{as required.}$$

Next we prove the result for T^S. Let $\Omega_K^\circ = \{j \in \{1, 2, \ldots, K\}: \beta_j \alpha_j = \tau^\circ\}$. Combining (4.2) and (4.3) we obtain

$$\lim_{x \to \infty} \frac{\log \mathbb{P}[\gamma_j L > x]}{\log \mathbb{P}[A^j > x]} = \lim_{x \to \infty} \frac{\beta_j \log \bar{F}(x)}{\log \bar{G}_j(x)} = \beta_j \alpha_j = \tau^\circ, \qquad j \in \Omega_K^\circ,$$

which, by Proposition 2.1, yields

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T_j > t]}{\log t} = -\tau^\circ, \qquad j \in \Omega_K^\circ.$$

Since $T^S = \max_{1 \le j \le K} T_j$, we have

$$\mathbb{P}[T^S > t] \ge \max_{j \in \Omega_K^\circ} \mathbb{P}[T_j > t],$$

and, hence,

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^S > t]}{\log t} \ge \max_{j \in \Omega_K^\circ} \lim_{t \to \infty} \frac{\log \mathbb{P}[T^S > t]}{\log t} = -\tau^\circ. \tag{4.5}$$

On the other hand, for $0 < \psi_j < 1/\mathbb{E}[A^j + U^j]$,

$$\mathbb{P}[T^S > t] \le \sum_{j=1}^K \mathbb{P}[T_j > t] \le K \max_{1 \le j \le K} \mathbb{P}[N_j > \psi_j t] + \sum_{j=1}^K \mathbb{P}[T_j > t, \, N_j \le \psi_j t].$$

Using (4.4),

$$\lim_{t \to \infty} \frac{\log(K \max_{1 \le j \le K} \mathbb{P}[N_j > \psi_j t])}{\log t} = \max_{1 \le j \le K} \lim_{t \to \infty} \frac{\log \mathbb{P}[N_j > \psi_j t]}{\log t} = -\tau^\circ,$$

so $\max_{1 \le j \le K} \mathbb{P}[N_j > \psi_j t] = t^{-\tau^\circ + o(1)}$. By Lemma 3.3, for some $\nu > \tau^\circ$, we have $\mathbb{P}[T_j > t, \, N_j \le \psi_j t] = O(t^{-\nu}) = o(\max_{1 \le j \le K} \mathbb{P}[N_j > \psi_j t])$. Therefore,

$$\lim_{t \to \infty} \frac{\log \mathbb{P}[T^S > t]}{\log t} \le \lim_{t \to \infty} \frac{\log(K \max_{1 \le j \le K} \mathbb{P}[N_j > \psi_j t])}{\log t} = -\tau^\circ,$$

which, combined with (4.5), completes the proof.

Theorem 4.2 characterizes the tail performance of the split transmission scheme in terms of the logarithmic asymptotics. Next, we present a theorem on the more refined asymptotic result.

Theorem 4.3. *Suppose that*

$$\bar{F}(x)^{-1} \sim \zeta_j \Phi_j(\bar{G}_j(x)^{-1}) \tag{4.6}$$

and

$$\bar{F}(x)^{-1} \sim \xi_j \Theta_j(\bar{F}(\gamma_j x)^{-1}), \tag{4.7}$$

where $\zeta_j, \xi_j > 0$, and $\Phi_j(\cdot), \Theta_j(\cdot)$ are regularly varying with indices $\alpha_j > 0, \beta_j > 0$, respectively, such that $\Theta_i(\Phi(x)) \sim \Theta_j(\Phi(x))$ if $\beta_i \alpha_i = \beta_j \alpha_j$. Let $\tau^\circ = \min_{1 \le j \le K} \beta_j \alpha_j$. If the moment conditions hold with parameter τ° then, as $t \to \infty$,

$$\Pi^\circ(t)\mathbb{P}[T^S > t] \to \sum_{\{J: \, \varnothing \ne J \subset \Omega_K^\circ\}} \frac{(-1)^{|J|+1} \Gamma(\tau^\circ + 1)}{(\sum_{j \in J} (\mathbb{E}[A^j + U^j])^{-1} \xi_j^{1/\tau^\circ} \zeta_j^{1/\alpha_j})^{\tau^\circ}}, \tag{4.8}$$

where $\Omega_K^\circ = \{j \in \{1, \ldots, K\}: \beta_j \alpha_j = \tau^\circ\}$ and $\Pi^\circ(t) \sim \Theta_j(\Phi_j(t))$ for $j \in \Omega_K^\circ$.

Note that (4.6) and (4.7) are strengthened versions of (4.2) and (4.3), respectively. If the limit in (4.8) is not 0, e.g. when $|\Omega_K^\circ| = 1$, then we obtain an asymptotic representation of $\mathbb{P}[T^S > t]$.

Proof of Theorem 4.3. By (4.6) and (4.7),

$$\bar{F}(x)^{-1} \sim \xi_j \Theta_j(\bar{F}(\gamma_j x)^{-1}) \sim \xi_j \zeta_j^{\beta_j} \Pi_j(\bar{G}_j(\gamma_j x)^{-1}),$$

where $\Pi_j := \Theta_j \circ \Phi_j$ is regularly varying with index $\tau_j := \beta_j \alpha_j > 0$. By Lemma 3.1,

$$\mathbb{P}\left[\bigcap_{j \in J}\{N_j > \psi_j t\}\right] \sim \frac{\Gamma(\tau_J^* + 1)}{(\sum_{j \in J^*} \psi_j \xi_j^{1/\tau_J^*} \zeta_j^{1/\alpha_j})^{\tau_J^*}} \frac{1}{\Pi_J^*(t)}, \tag{4.9}$$

where $\tau_J^* = \max_{j \in J} \tau_j$, $J^* = \{j \in J : \tau_j = \tau_J^*\}$, and $\Pi_J^*(t) \sim \Pi_j(t)$ for $j \in J^*$. By the inclusion–exclusion principle,

$$\mathbb{P}\left[\bigcup_{j=1}^K\{N_j > \psi_j t\}\right] = \sum_{\varnothing \neq J \subset \{1,2,...,K\}} (-1)^{|J|+1}\mathbb{P}\left[\bigcap_{j \in J}\{N_j > \psi_j t\}\right],$$

which, together with (4.9), yields, for any $\psi_j > 0$, as $t \to \infty$,

$$\Pi^\circ(t)\mathbb{P}\left[\bigcup_{j=1}^K\{N_j > \psi_j t\}\right] \to \sum_{\{J:\ \varnothing \neq J \subset \Omega_K^\circ\}} \frac{(-1)^{|J|+1}\Gamma(\tau^\circ + 1)}{(\sum_{j \in J} \psi_j \xi_j^{1/\tau^\circ} \zeta_j^{1/\alpha_j})^{\tau^\circ}}. \tag{4.10}$$

Now let $\widehat{\psi}_j < 1/\mathbb{E}[A^j + U^j] < \widetilde{\psi}_j$. By union bounds,

$$\begin{aligned}
\mathbb{P}&\left[\bigcup_{j=1}^K\{N_j > \widetilde{\psi}_j t\}\right] - \sum_{j=1}^K \mathbb{P}[T_j \leq t, N_j > \widetilde{\psi}_j t] \\
&\leq \mathbb{P}[T^S > t] \\
&= \mathbb{P}\left[\bigcup_{j=1}^K\{T_j > t\}\right] \\
&\leq \mathbb{P}\left[\bigcup_{j=1}^K\{N_j > \widehat{\psi}_j t\}\right] + \sum_{j=1}^K \mathbb{P}[T_j > t, N_j \leq \widehat{\psi}_j t].
\end{aligned}$$

By Lemma 3.2, Lemma 3.3, (4.10), and Proposition 1.5.1 of [4],

$$\sum_{\{J:\ \varnothing \neq J \subset \Omega_K^\circ\}} \frac{(-1)^{|J|+1}\Gamma(\tau^\circ + 1)}{(\sum_{j \in J} \widetilde{\psi}_j \xi_j^{1/\tau^\circ} \zeta_j^{1/\alpha_j})^{\tau^\circ}} \leq \lim_{t \to \infty} \Pi^\circ(t)\mathbb{P}[T^S > t]$$

$$\leq \sum_{\{J:\ \varnothing \neq J \subset \Omega_K^\circ\}} \frac{(-1)^{|J|+1}\Gamma(\tau^\circ + 1)}{(\sum_{j \in J} \widehat{\psi}_j \xi_j^{1/\tau^\circ} \zeta_j^{1/\alpha_j})^{\tau^\circ}}.$$

Now letting $\widehat{\psi}_j, \widetilde{\psi}_j \to 1/\mathbb{E}[A^j + U^j]$ completes the proof.

4.3. Optimal split transmission

According to Theorem 4.2, in order to minimize the power law tail exponent of the delay distribution, the γ_js should be chosen in such a way that $\min_{1 \leq j \leq K} \beta_j \alpha_j$ is maximized. We may speculate that we need to choose the γ_js so that $\beta_1 \alpha_1 = \beta_2 \alpha_2 = \cdots = \beta_K \alpha_K$. The following theorem confirms that this is indeed the case when $\log(\bar{F}(x)^{-1})$ is not slowly varying. A related work on optimal file split under a different problem setting can be found in [7].

Theorem 4.4. *Suppose that we use split transmission over K heterogeneous paths, each satisfying (4.2). If the limit*

$$\beta(\gamma) = \lim_{x \to \infty} \frac{\log \bar{F}(x)}{\log \bar{F}(\gamma x)}$$

exists for all $0 < \gamma < 1$ then there exists a unique constant $\rho \geq 0$ such $\beta(\gamma) = \gamma^{-\rho}$. Let

$$\alpha_\rho = \begin{cases} \left(\sum_{i=1}^{K} \alpha_i^{1/\rho} \right)^\rho, & \rho > 0, \\ \max_{1 \leq i \leq K} \alpha_i, & \rho = 0. \end{cases}$$

If, in addition, the moment conditions hold with parameter α_ρ then the minimum power law tail exponent achievable is $-\alpha_\rho$. The optimal splitting scheme that achieves the minimum is as follows.

(a) *If $\rho > 0$ then*

$$\gamma_j^* = \frac{\alpha_j^{1/\rho}}{\sum_{i=1}^{K} \alpha_i^{1/\rho}}, \qquad j = 1, 2, \ldots, K. \tag{4.11}$$

(b) *If $\rho = 0$ then $\gamma_j = 0$ for $\alpha_j \neq \max_{1 \leq i \leq K} \alpha_i$ and the other γ_j can be any partition of one.*

In the preceding result, our objective is to minimize the power law tail exponent. When $\rho = 0$, we have $\beta(\gamma) = 1$, and $\log \bar{F}(x)^{-1}$ is a slowly varying function. In this case, we should only use the best paths (i.e. paths with the largest α_j value), and the scheme in (4.11) is to split the packet arbitrarily among the best paths. This provides us with some unused degrees of freedom that may potentially be used to optimize some additional objectives, but we will not pursue this here. When $\rho > 0$, all the channels are utilized, and the optimal fraction over each path is specified by (4.11). In this case, one can easily check that the optimal tail exponent is indeed achieved when $\beta_1 \alpha_1 = \beta_2 \alpha_2 = \cdots = \beta_K \alpha_K$.

Note that $\alpha_\rho = (\sum_{i=1}^{K} \alpha_i^{1/\rho})^\rho \geq \alpha^*$ with equality if and only if $\rho = 0$, where $\alpha^* = \max_{1 \leq j \leq K} \alpha_j > 0$, as defined in Theorem 3.2. Thus, under the assumption of Theorem 4.4, split transmission achieves a better exponent than redundant transmission if $\rho > 0$.

Proof of Theorem 4.4. (a) Note that $\beta(\gamma) \geq 1$ on $(0, 1)$. If $\beta(\gamma) = 1$ for all $\gamma \in (0, 1)$ then $\beta(\gamma) = \gamma^{-\rho}$ for $\rho = 0$. Now assume that $\beta_0 = \beta(\gamma_0) > 1$ for some $\gamma_0 \in (0, 1)$. Observe that $\beta(\gamma_1 \gamma_2) = \beta(\gamma_1)\beta(\gamma_2)$ for any $\gamma_1, \gamma_2 \in (0, 1)$. Thus, for any positive integers m and n,

$$\beta(\gamma_0^{m/n}) = (\beta(\gamma_0^{1/n}))^{n \times m/n} = (\beta((\gamma_0^{1/n})^n))^{m/n} = \beta_0^{m/n}.$$

Since β is monotonically decreasing and the positive rationals are dense in $\mathbb{R}^+$,

$$\beta(\gamma_0^r) = \beta_0^r, \qquad r \in \mathbb{R}^+,$$

or, equivalently,

$$\beta(\gamma) = \gamma^{\log \beta_0 / \log \gamma_0} = \gamma^{-\rho}, \qquad \gamma \in (0, 1),$$

where $\rho = -\log \beta_0 / \log \gamma_0 > 0$. It is clear that ρ is unique.

(b) Let $\{\gamma_j\}$ be any splitting scheme. Let

$$\tau^\circ = \min_{\{j \,:\, \gamma_j > 0\}} \alpha_j \gamma_j^{-\rho}. \tag{4.12}$$

If $\rho = 0$ then

$$\tau^\circ = \min_{\{j \,:\, \gamma_j > 0\}} \alpha_j \le \max_{1 \le j \le K} \alpha_j = \alpha_\rho$$

with equality if and only if $\gamma_j = 0$ whenever $\alpha_j \neq \alpha_\rho$. If $\rho > 0$ then (4.12) gives

$$\gamma_j (\tau^\circ)^{1/\rho} \le \alpha_j^{1/\rho}, \qquad j = 1, 2, \ldots, K.$$

Summing over j and noting that $\sum_j \gamma_j = 1$, we have $(\tau^\circ)^{1/\rho} \le \sum_{j=1}^{K} \alpha_j^{1/\rho}$, or $\tau^\circ \le \alpha_\rho$, with equality if $\gamma_j = \gamma_j^*$ as given by (4.11). In both cases, Theorem 4.2 shows that the minimum power law tail exponent achievable is $-\max \tau^\circ = -\alpha_\rho$.

To illustrate the result of Theorem 4.4, we compute the optimal splitting scheme for some typical distributions.

Example 4.3. (*Weibull distribution.*) Consider the heterogeneous counterpart of Example 4.1. Suppose that the packet length L and all the available periods A^j ($1 \le j \le K$) follow Weibull distributions, i.e.

$$\bar{F}(x) = \mathbb{P}[L > x] = e^{-(\lambda x)^b}, \qquad \bar{G}_j(x) = \mathbb{P}[A^j > x] = e^{-(\mu_j x)^b},$$

where $\lambda > 0$, $\mu_j > 0$, and $b > 0$. Then $\alpha_j = (\lambda / \mu_j)^b$, $\beta(\gamma) = \gamma^{-b}$, and $\rho = b$. Therefore, the optimal splitting scheme is

$$\gamma_j = \frac{(\lambda / \mu_j)^{1/b}}{\sum_{i=1}^{K} (\lambda / \mu_i)^{1/b}} = \frac{\mu_j^{-1/b}}{\sum_{i=1}^{K} \mu_i^{-1/b}}, \qquad j = 1, \ldots, K.$$

Example 4.4. (*Pareto distribution.*) Suppose that the packet size L and all the available periods A^j follow Pareto distributions, i.e.

$$\bar{F}(x) = \mathbb{P}[L > x] = \begin{cases} \left(\dfrac{b_0}{x}\right)^\lambda, & x \ge b_0, \\ 1, & x < b_0, \end{cases}$$

$$\bar{G}_j(x) = \mathbb{P}[A^j > x] = \begin{cases} \left(\dfrac{b_j}{x}\right)^\mu, & x \ge b_j, \\ 1, & x < b_j. \end{cases}$$

As noted in Example 4.2, we have $\beta(\gamma) = 1$ and $\rho = 0$. Thus, the optimal splitting scheme is to use the best paths only, i.e. γ_j is nonzero only if $\alpha_j = \max_{1 \le i \le K} \alpha_i$, and the split among these paths is arbitrary as long as the tail exponent is concerned.

Appendix A. Proof of Lemma 3.1

The proof itself is divided into several lemmas. We first recall the following result.

Lemma A.1. (Proposition 1.5.8 of [4].) *Let $\Phi(x)$ be regularly varying with index $\alpha > 0$. For large enough x_0, the function defined by*

$$\Psi(x) = \int_{x_0}^{x} \alpha u^{-1} \Phi(u)\,du, \qquad x \geq x_0, \tag{A.1}$$

satisfies $\Psi(x) \sim \Phi(x)$.

The key step is the following lemma.

Lemma A.2. *Let $\Phi(x)$ be regularly varying with index $\alpha > 0$ and continuous on $[x_0, \infty)$ for some $x_0 > 0$. Let $\Psi(x)$ be given by (A.1), and let $\Psi^{\leftarrow}(x)$ be its inverse. Then, for all small enough $\varepsilon > 0$, as $x \to \infty$,*

$$H(x; \varepsilon, C) := \int_0^{\varepsilon} \exp\left(-\frac{Cx}{\Psi^{\leftarrow}(v^{-1})}\right) dv \sim \frac{\Gamma(\alpha + 1)}{C^\alpha} \frac{1}{\Phi(x)}. \tag{A.2}$$

Proof. Note that $\Psi(x)$ is a monotonically increasing diffeomorphism from $[x_0, \infty)$ onto $[0, \infty)$. Changing variable according to $u = x/\Psi^{\leftarrow}(v^{-1})$, i.e. $v = 1/\Psi(x/u)$,

$$H(x; \varepsilon, C) = \int_0^{c(\varepsilon)x} \alpha u^{-1} e^{-Cu} \frac{\Phi(x/u)}{\Psi^2(x/u)}\,du, \tag{A.3}$$

where $c(\varepsilon) = 1/\Psi^{\leftarrow}(\varepsilon^{-1}) > 0$.

Note that $\Phi^*(x) = x^{\alpha/2}/\Phi(x)$ is regularly varying with index $-\alpha/2 < 0$. An application of Theorem 1.5.2 of [4] to $\Phi^*(x)$ implies that there exists M_0 such that, for $x > M_0$ and $0 < u \leq 1$,

$$u^{-\alpha/2} \frac{\Phi(x)}{\Phi(x/u)} = \frac{\Phi^*(x/u)}{\Phi^*(x)} \leq u^{\alpha/2} + 1 \leq 2,$$

and, hence,

$$\frac{\Phi(x)}{\Phi(x/u)} \leq 2u^{\alpha/2}. \tag{A.4}$$

By Theorem 1.5.6 of [4], there exists $M_1 > M_0$ such that, for $x \geq M_1$ and $1 \leq u \leq x/M_1$,

$$\frac{\Phi(x)}{\Phi(x/u)} \leq 2u^{\alpha+1}. \tag{A.5}$$

Since $\Psi(x) \sim \Phi(x)$, there exists $M \geq M_1$ such that, for all $x/u \geq M$,

$$\frac{\Phi(x/u)}{\Psi(x/u)} \leq 2. \tag{A.6}$$

Since $c(\varepsilon) = 1/\Psi^{\leftarrow}(\varepsilon^{-1}) \to 0$ as $\varepsilon \to 0$, there exists $\varepsilon_0 > 0$ such that $c(\varepsilon) < 1/M$ for all $\varepsilon < \varepsilon_0$. Combining (A.4), (A.5), and (A.6) yields

$$\frac{\Phi(x)\Phi(x/u)}{\Psi^2(x/u)} = \frac{\Phi(x)}{\Phi(x/u)}\left(\frac{\Phi(x/u)}{\Psi(x/u)}\right)^2 \leq 8u^{\alpha+1} + 8u^{\alpha/2} \tag{A.7}$$

for $x \geq M$, $\varepsilon < \varepsilon_0$, and $0 < u \leq c(\varepsilon)x$.

Note that, as $x \to \infty$,

$$f(u, x) := \alpha u^{-1} \mathrm{e}^{-Cu} \frac{\Phi(x)\Phi(x/u)}{\Psi^2(x/u)} \mathbf{1}(0 < u \le c(\varepsilon)x] \to \alpha u^{\alpha-1} \mathrm{e}^{-Cu}.$$

Moreover, for $x \ge M$ and $\varepsilon < \varepsilon_0$, (A.7) yields

$$0 \le f(u, x) \le g(u) := 8\alpha(u^{\alpha} + u^{\alpha/2-1})\mathrm{e}^{-Cu},$$

where $g(u) \in L^1(0, \infty)$ with integral

$$\int_0^{\infty} g(u)\,\mathrm{d}u = \frac{8\alpha\Gamma(\alpha+1)}{C^{\alpha+1}} + \frac{8\alpha\Gamma(\alpha/2)}{C^{\alpha/2}} < \infty.$$

Therefore, by the dominated convergence theorem and (A.3), for $\varepsilon < \varepsilon_0$,

$$\Phi(x)H(x; \varepsilon, C) = \int_0^{\infty} f(u, x)\,\mathrm{d}u \to \int_0^{\infty} \alpha\mathrm{e}^{-Cu}u^{\alpha-1}\,\mathrm{d}u = \frac{\Gamma(\alpha+1)}{C^{\alpha}}$$

as $x \to \infty$.

Lemma A.3. *Let $\Phi(x)$ be regularly varying with index $\alpha > 0$ and continuous on $[x_0, \infty)$ for some $x_0 > 0$. Let $\Psi(x)$ be given by (A.1), and let $\Psi^{\leftarrow}(x)$ be its inverse. If $h(x) \sim C/\Psi^{\leftarrow}(\bar{F}(x)^{-1})$ then, for all large enough z,*

$$\mathbb{E}[\mathrm{e}^{-th(L)}\mathbf{1}(L > z)] \sim \frac{\Gamma(\alpha+1)}{C^{\alpha}}\frac{1}{\Phi(t)} \quad \text{as } t \to \infty.$$

Proof. Given $\delta \in (0, 1)$, for all large enough x,

$$\frac{(1-\delta)C}{\Psi^{\leftarrow}(\bar{F}(x)^{-1})} \le h(x) \le \frac{(1+\delta)C}{\Psi^{\leftarrow}(\bar{F}(x)^{-1})}.$$

Thus, for all large enough z, after integrating and changing variables according to $v = \bar{F}(x)$, we obtain

$$H(t; \bar{F}(z), (1+\delta)C) \le \mathbb{E}[\mathrm{e}^{-th(L)}\mathbf{1}(L > z)] \le H(t; \bar{F}(z), (1-\delta)C),$$

where $H(t; \varepsilon, C)$ is as defined in (A.2). When z is large enough, $\bar{F}(z)$ is small enough, so, by (A.2),

$$\frac{\Gamma(\alpha+1)}{(1+\delta)^{\alpha}C^{\alpha}} \le \lim_{t \to \infty} \Phi(t)\mathbb{E}[\mathrm{e}^{-th(L)}\mathbf{1}(L > z)] \le \frac{\Gamma(\alpha+1)}{(1-\delta)^{\alpha}C^{\alpha}}.$$

Now letting $\delta \to 0$ yields the desired result (A.8).

Lemma A.4. *Let $\Phi(x)$ be regularly varying with index $\alpha > 0$. Let $f(x)$ and $g(x)$ tend to ∞ as $x \to \infty$. If $f(x) \sim g(x)$ then $\Phi(f(x)) \sim \Phi(g(x))$.*

Proof. Let $\Psi(x)$ be given by (A.1). Given any $\varepsilon \in (0, 1)$, for all large enough x,

$$(1-\varepsilon)g(x) \le f(x) \le (1+\varepsilon)g(x),$$

and, hence, the mononicity of $\Psi(x)$ yields

$$\Psi((1-\varepsilon)g(x)) \le \Psi(f(x)) \le \Psi((1+\varepsilon)g(x)).$$

Since $\Psi(x)$ is regularly varying with index α, by (2.1),

$$
\begin{aligned}
(1 - \varepsilon)^\alpha &= \lim_{x \to \infty} \frac{\Psi((1 - \varepsilon)g(x))}{\Psi(g(x))} \\
&\leq \varliminf_{x \to \infty} \frac{\Psi(f(x))}{\Psi(g(x))} \\
&\leq \varlimsup_{x \to \infty} \frac{\Psi(f(x))}{\Psi(g(x))} \\
&\leq \lim_{x \to \infty} \frac{\Psi((1 + \varepsilon)g(x))}{\Psi(g(x))} \\
&= (1 + \varepsilon)^\alpha.
\end{aligned}
$$

Letting $\varepsilon \to 0$ yields $\Psi(f(x)) \sim \Psi(g(x))$. Since $\Phi(x) \sim \Psi(x)$, it follows that $\Phi(f(x)) \sim \Phi(g(x))$.

Proof of Lemma 3.1. Replacing $\Phi_j(x)$ by $\Psi_j(x)$ as given in (A.1) if necessary, we can assume that $\Phi_j(x)$ is continuous on $[x_0, \infty)$ for some large enough x_0. Now let $\Psi_j(x)$ be given by (A.1), and let $\Psi_j^{\leftarrow}(x)$ be its inverse. By Theorem 1.5.12 of [4], $\Psi_j^{\leftarrow}(x)$ is regularly varying with index $1/\alpha_j$. Using $\Phi_j(x) \sim \Psi_j(x)$, we obtain

$$
\Psi_j(\bar{G}_j(\gamma_j x)^{-1}) \sim \Phi_j(\bar{G}_j(\gamma_j x)^{-1}) \sim \zeta_j^{-1} \bar{F}(x)^{-1},
$$

which, by Lemma A.4, yields

$$
\bar{G}_j(\gamma_j x)^{-1} \sim \Psi_j^{\leftarrow}(\zeta_j^{-1} \bar{F}(x)^{-1}) \sim \zeta_j^{-1/\alpha_j} \Psi_j^{\leftarrow}(\bar{F}(x)^{-1}),
$$

and, hence,

$$
\sum_{j \in J} \psi_j \bar{G}_j(\gamma_j x) \sim \sum_{j \in J} \frac{\psi_j \zeta_j^{1/\alpha_j}}{\Psi_j^{\leftarrow}(\bar{F}(x)^{-1})} \sim \sum_{j \in J^*} \frac{\psi_j \zeta_j^{1/\alpha_j^*}}{\Psi_{J^*}^{\leftarrow}(\bar{F}(x)^{-1})},
$$

where $\alpha_J^* = \max_{j \in J} \alpha_j$, $J^* = \{j \in J : \alpha_j = \alpha_J^*\}$, and $\Psi_{J^*}^{\leftarrow}(x)$ is the inverse of $\Psi_J^*(x)$, which corresponds to $\Phi_J^*(x)$ as in (A.1). Thus, by Lemma A.3, for all large enough z,

$$
Q(t, z) := \mathbb{E}\left[\exp\left(-t \sum_{j \in J} \psi_j \bar{G}_j(L)\right) \mathbf{1}(L > z)\right] \sim \frac{\Gamma(\alpha_J^* + 1)}{\left(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*}\right)^{\alpha_J^*}} \frac{1}{\Phi_J^*(t)}. \quad \text{(A.8)}
$$

Denote the left-hand side of (3.5) by $R(t)$. Since the N_js are independent conditioned on L,

$$
R(t) = \mathbb{E}\left[\prod_{j \in J} \mathbb{P}[N_j > \psi_j t \mid L]\right] = \mathbb{E}\left[\prod_{j \in J} (1 - \bar{G}_j(\gamma_j L))^{\lfloor \psi_j t \rfloor}\right]. \quad \text{(A.9)}
$$

Note that, given any $\varepsilon > 0$, there exists $M > 0$ such that, for all $x > M$,

$$
\prod_{j \in J} (1 - \bar{G}_j(\gamma_j x))^{\lfloor \psi_j t \rfloor} \geq \prod_{j \in J} (1 - \bar{G}_j(\gamma_j x))^{\psi_j t} \geq (1 - \varepsilon) \exp\left(-t \sum_{j \in J} \psi_j \bar{G}_j(\gamma_j x)\right).
$$

Thus, for all large enough z,

$$R(t) \geq \mathbb{E}\left[\prod_{j \in J}(1 - \bar{G}_j(\gamma_j L))^{\lfloor \psi_j t \rfloor} \mathbf{1}(L > z)\right] \geq (1 - \varepsilon)Q(t, z),$$

which together with (A.8) yields

$$\lim_{t \to \infty} \Phi_J^*(t)R(t) \geq (1 - \varepsilon)\frac{\Gamma(\alpha_J^* + 1)}{(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*})^{\alpha_J^*}}.$$

Letting $\varepsilon \to 0$,

$$\lim_{t \to \infty} \Phi_J^*(t)R(t) \geq \frac{\Gamma(\alpha_J^* + 1)}{(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*})^{\alpha_J^*}}. \tag{A.10}$$

On the other hand, the inequalities $\lfloor x \rfloor \geq x - 1$ and $1 - x \leq e^{-x}$ yield

$$\prod_{j \in J}(1 - \bar{G}_j(\gamma_j x))^{\lfloor \psi_j t \rfloor} \leq \exp\left(\sum_{j \in J}\bar{G}_j(\gamma_j x)\right)\exp\left(-t\sum_{j \in J}\psi_j \bar{G}_j(\gamma_j x)\right),$$

whence, by splitting (A.9) into two parts according to $L > z$ and $L \leq z$,

$$R(t) \leq \exp\left(\sum_{j \in J}\bar{G}_j(\gamma_j z)\right)Q(t, z) + \exp\left(-t\sum_{j \in J}\psi_j \bar{G}_j(\gamma_j z) + |J|\right). \tag{A.11}$$

By Proposition 1.5.1 of [4], the last term of (A.11) is $o(t^{-\alpha_J^* - 1}) = o(1/\Phi_J^*(t))$ as $t \to \infty$. Using (A.8), we obtain, for all large enough z,

$$\overline{\lim}_{t \to \infty} \Phi_J^*(t)R(t) \leq \exp\left(\sum_{j \in J}\bar{G}_j(z)\right)\frac{\Gamma(\alpha_J^* + 1)}{(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*})^{\alpha_J^*}}.$$

Now letting $z \to \infty$,

$$\overline{\lim}_{t \to \infty} \Phi(t)R(t) \leq \frac{\Gamma(\alpha_J^* + 1)}{(\sum_{j \in J^*} \psi_j \zeta_j^{1/\alpha_J^*})^{\alpha_J^*}}.$$

which together with (A.10) yields (3.5).

Appendix B. Proof of Lemma 3.3

The proof is divided into several steps. We first recall the following two results from [13].

Lemma B.1. (Corollary 1.6 of [13].) *Let $X_1, X_2, \ldots, X_n$ be i.i.d. random variables $\sim X$ such that $\mathbb{E}X = 0$ and $a_s^+ := \mathbb{E}[(X \vee 0)^s] < \infty$ for $1 \leq s \leq 2$. Then, for $x > y > (4na_s^+)^{1/s}$,*

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

Lemma B.2. (Corollary 1.8 of [13].) *Let $X_1, X_2, \ldots, X_n$ be i.i.d. random variables $\sim X$ such that $\mathbb{E}X = 0$, $\sigma^2 = \mathrm{var}[X] < \infty$, and $a_s^+ := \mathbb{E}[X^s \mathbf{1}(X \geq 0)] < \infty$ for $s \geq 2$. Then*

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i \geq x\right] \leq \frac{c_s a_s^+ n}{x^s} + \exp\left(-\frac{d_s x^2}{\sigma^2 n}\right),$$

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

We will use the two lemmas in the following combined form.

Corollary B.1. *Let $X_1, X_2, \ldots, X_n$ be i.i.d. random variables $\sim X$ such that $\mathbb{E}X = 0$ and $\mathbb{E}[X^s] < \infty$ for some $s \geq 1$. If $n = O(x^q)$ for some $q < s \wedge 2$ then*

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i > x\right] = O\left(\frac{n}{x^s}\right) \quad \text{as } x \to \infty.$$

Proof. If $1 \leq s \leq 2$ then $(4na_s^+)^{1/s} = O(x^{q/s}) = o(x)$, so $x/2 > (4na_s^+)^{1/s}$ for large enough x. Setting $y = x/2$ in Lemma B.1 and then applying Markov's inequality to $\mathbb{P}[X > x/2]$ yields

$$\mathbb{P}\left[\sum_{i=1}^{n} X_i > x\right] \leq n\mathbb{P}\left[X > \frac{x}{2}\right] + n\frac{2^{s-1}e^2 a_s^+}{x^s} \leq 2^{s-1}(2 + e^2)a_s^+ \frac{n}{x^s}.$$

If $s \geq 2$ then $x^2/n = \Omega(x^{2-q})$ and the result follows from Lemma B.2. ∎

The next two lemmas are the key ingredients for the proof of Lemma 3.3.

Lemma B.3. *If $\mathbb{E}[(U^j)^s] < \infty$ for some $s > \alpha \vee 1$ then there exists an $\nu > \alpha$ such that, as $t \to \infty$,*

$$\mathbb{P}\left[\sum_{i=1}^{N_j}(U_i^j - \mathbb{E}[U^j]) > \delta t, \; N_j \leq \psi_j t\right] = O\left(\frac{1}{t^\nu}\right).$$

Proof. Let $\tilde{U}_i^j = U_i^j - \mathbb{E}[U^j]$. Since N_j and $\{U_i^j\}$ are independent,

$$\mathbb{P}\left[\sum_{i=1}^{N_j} \tilde{U}_i^j > \delta t, \; N_j \leq \psi_j t\right] = \sum_{n=1}^{M} \mathbb{P}[N_j = n]\mathbb{P}\left[\sum_{i=1}^{n} \tilde{U}_i^j > \delta t\right],$$

where $M = \lfloor \psi_j t \rfloor$. By Corollary B.1, the right-hand side is

$$\sum_{n=1}^{M} \mathbb{P}[N_j = n]O(nt^{-s}) = O\left(t^{-s} \sum_{n=1}^{M} n\mathbb{P}[N_j = n]\right).$$

Using summation by parts,

$$\sum_{n=1}^{M} n\mathbb{P}[N_j = n] = 1 + \sum_{n=1}^{M-1} \mathbb{P}[N_j > n] - M\mathbb{P}[N_j > M] \leq 2 + \sum_{n=2}^{M} \mathbb{P}[N_j > n].$$

If $\alpha > 1$ then let $\theta \in (1, \alpha)$; otherwise, let $\theta \in (1 + \alpha - s, \alpha)$. By Lemma 3.1, there exists a constant D_θ such that $\mathbb{P}[N_j > n] \leq D_\theta n^{-\theta}$. Thus,

$$\sum_{n=2}^{M} \mathbb{P}[N_j > n] \leq \sum_{n=2}^{M} \frac{D_\theta}{n^\theta} \leq \int_1^M \frac{D_\theta}{x^\theta} \, \mathrm{d}x = \frac{D_\theta}{1-\theta}[M^{1-\theta} - 1] = O(t^{(1-\theta)\vee 0})$$

and

$$\mathbb{P}\left[\sum_{i=1}^{N_j} \tilde{U}_i^j > \delta t, \ N_j \leq \psi_j t\right] = O(t^{-s} t^{(1-\theta)\vee 0}) = O(t^{-\nu}),$$

where $\nu = s \wedge (s + \theta - 1) > \alpha$.

Lemma B.4. *Let X and Y be positive random variables such that $\mathbb{E}[X^{1+\theta}] < \infty$ for some $\theta > 0$, and $\mathbb{E}[Y^s] < \infty$ for some $s > 0$. Let $\{X_i\}$ be i.i.d. $\sim X$. Then, for any $\psi < 1/\mathbb{E}X$ and $\delta < 1 - \psi \mathbb{E}X$,*

$$\mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} X_i \wedge Y > (1-\delta)t\right] = O\left(\frac{1}{t^s}\right).$$

Proof. Choose B such that $\mathbb{E}X < B < (1-\delta)/\psi$. Let $\eta = 1 - \delta - B\psi > 0$. Let $\{Z_i\}$ be i.i.d. exponential random variables $\sim Z$ that are independent of $\{X_i\}$, such that $\mathbb{E}X < \mathbb{E}Z < B$.

By Proposition X.1.1 of [2], $\sup_n \sum_{i=1}^n (Z_i - B)$ is equal in distribution to the steady-state waiting time of a D/M/1 queue with interarrival time D and service time Z. Theorem VIII.5.8 of [2] then yields

$$\mathbb{P}\left[\sup_n \sum_{i=1}^n (Z_i - B) > \frac{1}{2}\eta t\right] = o(t^{-s}).$$

By Proposition X.1.1 and Theorem VIII.5.7 of [2], $\sup_n \sum_{i=1}^n (X_i \wedge (\varepsilon t) - Z_i)$ is equal in distribution to the steady-state workload of an M/G/1 queue with interarrival time Z and truncated service time $X \wedge (\varepsilon t)$. By Lemma 3.2 of [8], there exists $\varepsilon > 0$ such that

$$\mathbb{P}\left[\sup_n \sum_{i=1}^n (X_i \wedge (\varepsilon t) - Z_i) > \frac{1}{2}\eta t\right] = o(t^{-s}).$$

Therefore,

$$\mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} X_i \wedge (\varepsilon t) > (1-\delta)t\right]$$

$$\leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} (X_i \wedge (\varepsilon t) - B) > \eta t\right]$$

$$\leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} (X_i \wedge (\varepsilon t) - Z_i) > \frac{1}{2}\eta t\right] + \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} (Z_i - B_i) > \frac{1}{2}\eta t\right]$$

$$= o(t^{-s}).$$

By Markov's inequality, $\mathbb{P}[Y > \varepsilon t] \leq \mathbb{E}[Y^s]/(\varepsilon t)^s = O(1/t^s)$. Thus,

$$\mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} X_i \wedge Y > (1-\delta)t\right] \leq \mathbb{P}\left[\sum_{i=1}^{\lfloor \psi t \rfloor} X_i \wedge (\varepsilon t) > (1-\delta)t\right] + \mathbb{P}[Y > \varepsilon t] = O(t^{-s}).$$

This completes the proof.

Proof of Lemma 3.3. Note that, for $N_j \leq \psi_j t$,

$$
\begin{aligned}
T_j &= \sum_{i=1}^{N_j-1} (A_i^j + U_i^j) + L_j \\
&\leq \sum_{i=1}^{N_j} (A_i^j \wedge L_j + U_i^j) \\
&= \sum_{i=1}^{N_j} (A_i^j \wedge L_j + \mathbb{E}[U_i^j]) + \sum_{i=1}^{N_j} (U_i^j - \mathbb{E}[U^j]) \\
&\leq \sum_{i=1}^{\lfloor \psi_j t \rfloor} (A_i^j \wedge L_j + \mathbb{E}[U_i^j]) + \sum_{i=1}^{N_j} (U_i^j - \mathbb{E}[U^j]).
\end{aligned}
$$

Thus,

$$
\begin{aligned}
\mathbb{P}[T_j > t, \, N_j \leq \psi_j t] \leq \; &\mathbb{P}\left[\sum_{i=1}^{\lfloor \psi_j t \rfloor} (A_i^j \wedge L_j + \mathbb{E}[U^j]) > (1-\delta)t \right] \\
&+ \mathbb{P}\left[\sum_{i=1}^{N_j} (U_i^j - \mathbb{E}[U^j]) > \delta t, \, N_j \leq \psi_j t \right].
\end{aligned}
$$

For $0 < \delta < 1 - \psi_j \mathbb{E}[A^j + U^j]$, the right-hand side is $O(t^{-\nu})$ for some $\nu > \alpha$ by Lemma B.3 and Lemma B.4. Note that the identity $A_i^j \wedge L_j + \mathbb{E}[U^j] = (A_i^j + \mathbb{E}[U^j]) \wedge (L_j + \mathbb{E}[U^j])$ has been used in the application of Lemma B.4.

Acknowledgements

References

[1] ANDERSEN, L. N. AND ASMUSSEN, S. (2008). Parallel computing, failure recovery, and extreme values. *J. Statist. Theory Appl.* **2,** 279–292.

[2] ASMUSSEN, S. (2003). *Applied Probability and Queues*, 2nd edn. Springer, New York.

[3] ASMUSSEN, S. *et al.* (2008). Asymptotic behavior of total times for jobs that must start over if a failure occurs. *Math. Operat. Res.* **33,** 932–944.

[4] BINGHAM, N. H., GOLDIE, C. M. AND TEUGELS, J. L. (1987). *Regular Variation* (Encyclopedia Math. Appl. **27**). Cambridge University Press.

[5] CANTOR, D. G. AND GERLA, M. (1974). Optimal routing in a packet-switched computer network. *IEEE Trans. Comput.* **23,** 1062–1069.

[6] GALLAGER, R. G. (1977). A minimum delay routing algorithm using distributed computation. *IEEE Trans. Commun.* **25,** 73–85.

[7] HOEKSTRA, G., MEI, R., NAZARATHY, Y. AND ZWART, B. (2009). Optimal file splitting for wireless networks with concurrent access. In *NET-COOP '09: Proc. 3rd Euro-NF Conf. on Network Control and Optimization* (Lecture Notes Comput. Sci. **5894**), Springer, Berlin, pp. 189–203.

[8] JELENKOVIĆ, P. AND MOMČILOVIĆ, P. (2003). Large deviation analysis of subexponential waiting times in a processor-sharing queue. *Math. Operat. Res.* **28,** 587–608.

[9] JELENKOVIĆ, P. R. AND TAN, J. (2007). Is ALOHA causing power law delays? In *Proc. 20th Internat. Teletraffic Congress* (Lecture Notes Comput. Sci. **4516**), Springer, Berlin, pp. 1149–1160.

[10] JELENKOVIĆ, P. R. AND TAN, J. (2007). Can retransmissions of superexponential documents cause subexponential delays? In *Proc. IEEE INFOCOM '07*, pp. 892–900.

[11] JELENKOVIĆ, P. R. AND TAN, J. (2007). Characterizing heavy-tailed distributions induced by retransmissions. Tech. Rep. EE2007-09-07, Department of Electrical Engineering, Columbia University. Available at http:// arxiv.org/abs/0709.1138v2.

[12] KLEINROCK, L. (1964). *Communication Nets: Stocastic Message Flow and Delay*. McGraw-Hill, New York.

[13] NAGAEV, S. V. (1979). Large deviations of sums of independent random variables. *Ann. Prob.* **7,** 745–789.

[14] SHEAHAN, R., LIPSKY, L., FIORINI, P. AND ASMUSSEN, S. (2006). On the completion time distribution for tasks that must restart from the beginning if a failure occurs. In *ACM SIGMETRICS Performance Evaluation Review*, Association for Computing Machinery, New York, pp. 24–26.

[15] TAN, J. AND SHROFF, N. (2010). Transition from heavy to light tails in retransmission durations. In *Proc. INFOCOM '10* (San Diego, California), IEEE Press, Piscataway, NJ, pp. 1334–1342.

[16] TAN, J., YANG, Y., SHROFF, N. B. AND GAMAL, H. E. (2011). Delay asymptotics with retransmissions and fixed rate codes over erasure channels. In *Proc. INFOCOM '11*, IEEE Press, Piscataway, NJ, pp. 1260–1268.