Can Multipath Mitigate Power Law Delays? – Effects of Parallelism on Tail Performance

Jian Tan[1], Wei Wei[2], Bo Jiang[2], Ness Shroff[1], Don Towsley[2]
[1] Department of Electrical Engineering, The Ohio State University, Columbus, OH 43210
[2]Department of Computer Science, University of Massachusetts, Amherst, MA 01003
UMass Computer Science Technical Report UM-CS-2009-055

Abstract

Parallelism has often been used to improve the reliability and efficiency of a variety of different engineering systems. In this paper, we quantify the efficiency of parallelism in systems that are prone to failures and exhibit power law transmission delays. We focus on the context of transmitting a data unit in communication networks, where parallelism can be achieved by multipath transmission (e.g., multipath routing). We investigate two types of transmission schemes: redundant and split transmission techniques. We find that the power-law transmission delay phenomenon still persists with multipath transmission. In particular, we show that when the transmission delays of each path are characterized by the same power law, redundant multipath transmission can only result in a constant factor performance gain, while order gains are possible when the delays are light tailed. We further compare the performance of redundant transmission and split transmission, and show that there is no clear winner. Depending on the packet size distribution properties and the manner in which splitting is performed, one scheme results in greater performance over the other. Specifically, split transmission 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. Based on our analysis, we develop an optimal split transmission strategy, and show that this strategy always outperforms redundant transmission.

I. Introduction

Parallelism is a common approach to improve reliability and efficiency in practice. For instance, in peer to peer systems, a file is downloaded in parallel from multiple peers; in grid computing, a job is allocated to multiple machines to be computed simultaneously; and in computer communication networks, multipath routing can be used to improve the efficiency and reliability of data transfer. In one type of parallelism, a file/job is fetched/computed in its entirety, and hence the completion time is the minimum of the completion times from/at the multiple locations. In another type of parallelism, a file/job is split into multiple pieces, fetched/computed independently, and hence the completion time is the maximum of the completion times of all the pieces. In both cases, we expect better efficiency from using parallelism since the delay is either the minimum one or because a smaller job needs to be completed.

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 retransmitted after a failure occurs. For example, in wireless communication networks, recent studies [8], [7], [9], [6] show that, contrary to traditional wisdom, when the probability of packet errors is a function of the packet length, retransmission-based protocols may cause power law transmission durations and possibly even zero throughput. Similar results have been reported in other contexts [12], [2]. A natural question to ask is whether and, if so, how, using parallelism can mitigate power law delays, which is the focus of our study.

To focus our discussion, let us consider the notion of parallelism in the context of communication networks, where a data unit can be transmitted using multiple paths (also known as multipath routing or more generally multipath transmission). A data unit can be a file or packet (which are used interchangeably, henceforward), and the transmission needs to restart after a failure (i.e., there is no check point in the transmission). We consider two multipath transmission strategies, *redundant* and *split transmission*, that correspond respectively to the two aforementioned types of parallelism. More specifically, redundant transmission replicates a packet and sends each copy over a different path and therefore, the transmission is successful once the first of the packets arrives at the destination; split transmission, on the other hand, breaks the data unit into several pieces and dispatches each piece along a different path, which completes the transmission when all the pieces arrive at the destination successfully.

We aim to answer the following three questions: (I) Can redundant or split transmission eliminate power laws in transmission delays, and how can the performance gain from multipath transmission be characterized? (II) Is split transmission or redundant transmission more beneficial in mitigating power law delays? and (III) What is the optimal strategy to split packets and dispatch those fragmented pieces to the appropriate paths.

To address the above questions, we generalize the single *channel* model introduced in [8] to a multipath channel model. First note that a channel can be viewed as a medium over which faults can occur causing jobs to be interrupted and retransmitted. In the context of communication networks, this corresponds to a wireless communication channel as in [8], in the context of grid computing the channel may correspond to the processor over which the computations are completed, etc. Henceforth, we will focus on communication networks and consider the notion of a channel in that context. Specifically, consider a communication network where there are K paths between a source and 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 period, A_i^j, and unavailable period, U_i^j. Only in each time period A_i^j when the channel becomes available, can a packet start its transmission over the path. If the length of A_i^j is longer than the length of the packet, the transmission is considered successful over path j; otherwise, we wait until the beginning of the next available period A_{i+1}^j and retransmit the packet from the beginning. The above model can be viewed as a first order approximation to channels that may fail. Channel failures can happen due to many reasons. For instance, in a wireless network environment, failures occur due to channel fading, interference and contention with other nodes, multipath effects, obstructions, and node mobility [11]. As a consequence, the signal to noise ratio (SINR) may vary in different time scales. The on periods $\{A_i^j\}$ in our model correspond to the situation when SINR is high, while the off periods $\{U_i^j\}$ correspond to the situation when SINR is low.

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

- We show that, when all packets are of the same size, redundant transmission can greatly reduce the transmission delay in the sense that the ratio of the delay distribution tail with and without redundant transmission tends to zero (see Proposition III.1). However, in reality, packet sizes are usually variable due to many other considerations, e.g., reducing communication costs and extra overhead induced from encapsulation. We prove that, when packet sizes are random variables that satisfy $\log \mathbb{P}[L > x] \approx \alpha^* \log \mathbb{P}[A^j > x]$, redundant transmission does not change the order of the probability tail of the transmission delays (see Theorem 2), and can only improve the system performance by a constant factor (see Theorem 3).

- We show that split transmission is effective in mitigating power 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 (see Theorems 4 and 5). To illustrate the point, we calculate the effectiveness of split transmission for different packet size distributions. Furthermore, we provide a solution for optimal split when we have heterogeneous paths, and show that this optimal strategy always outperforms redundant transmission (see Theorem 6). To refine the result, we also derive an exact asymptotic for packet delivery time under optimal split transmission (see Theorem 7).

In terms of related work, it was observed in [12] that power law processing times can arise in a system where jobs need to restart once a failure occurs. This observation was rigorously addressed in [8], [2], [7], [9] for a single channel model. The result reveals that, when the probability of packet errors is a function of the packet length, retransmission-based protocols could cause heavy-tailed (specifically, power law) transmission durations, even when the data units and channel characteristics are light-tailed. Our study generalizes the single channel model to the one with multiple paths. Multipath transmissions have also been studied in [1] using Extreme Value theory, but only when the number of paths goes to infinity. In this work, we focus on the context of multipath transmissions in computer networks with a fixed (possibly small) number of paths, where multipath transmission has long been used to improve reliability and efficiency (e.g., [10], [4], [5]).

Note that the specific investigation conducted in this paper has been in the context of data transmission in wireless communication networks. However, the mathematical setting described in Section 2 is quite general, and the results can be extended to many other applications 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. Section 2 presents the model description and some results on single path transmission. Redundant transmission and split transmission are investigated in Sections 3 and 4, respectively. Finally, Section 5 concludes the paper.

II. Model description and preliminary results

Let L be a random variable that denotes the length of a packet. Assume that there are $K \geq 1$ paths between the source and destination, as shown in Figure 1. 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 A_i^j and unavailable U_i^j periods, respectively.

Packet transmission can only be initiated at the start of an available period. For a packet transmission started at the beginning of A_i^j, if $A_i^j > L$, the transmission is considered successful over path j; otherwise, we wait until the beginning of the next available period A_{i+1}^j and retransmit the packet from the beginning.

We study two multipath transmission schemes, namely, redundant transmission and split transmission. Under redundant transmission, the same packet is transmitted over all K paths, and the transmission is successful as soon as one of the K duplicates arrives at the destination. Split transmission represents the strategy where a packet is split into

K pieces and each piece is sent over a different path. The transmission is complete once all the K pieces arrive at the destination successfully.

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

$$N_j \triangleq \inf\{i : A_i^j > L\},$$

and, the corresponding transmission time over this path is defined as

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

- Redundant transmission: the transmission completes when the first packet is successfully transmitted over one of the K paths. Therefore, the total transmission time T_r for this scheme satisfies

$$T_r \triangleq \min_{1 \leq j \leq K} T_j.$$

- Split transmission: the transmission completes when all K pieces of the packet are successfully transmitted. Therefore, the total transmission time T_s for this scheme satisfies

$$T_s \triangleq \max_{1 \leq j \leq K} T_j,$$

  and the total number of retransmissions over K paths is

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

In this paper, we assume that $\{U^j, U_i^j\}_{j \geq 1}$ and $\{A^j, A_i^j\}_{j \geq 1}$, $1 \leq j \leq K$ are mutually independent i.i.d. sequences of random variables, which are also independent of the packet size L. A sketch of the model depicting the system is shown in Figure 1.

We use the following notation to denote the complementary cumulative distribution functions for A^j, $1 \leq j \leq K$ and L,

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

and

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

We say K paths are *homogeneous* if $A^j \stackrel{d}{=} A$ and $U^j \stackrel{d}{=} U$ for $1 \leq j \leq K$, where "$\stackrel{d}{=}$" denotes equal in distribution. Accordingly, we use $\bar{G}(x) \triangleq \mathbb{P}[A > x]$. In general, $\{A^j\}_{1 \leq j \leq K}$ (and $\{U^j\}_{1 \leq j \leq K}$) need not be identically distributed, which represents the case of *heterogenous* paths.

Throughout this paper, a positive measurable function f is called regularly varying (at infinity) with index ρ if

$$\lim_{x \to \infty} f(\lambda x)/f(x) = \lambda^{\rho}$$

for all $\lambda > 0$. It is called slowly varying if $\rho = 0$ [3]. Additionally, for any two real functions $f(t)$ and $g(t)$, we use $f(t) \sim g(t)$ to denote $\lim_{t\to\infty} f(t)/g(t) = 1$. Similarly, we say that $f(t) \gtrsim g(t)$ if $\underline{\lim}_{t\to\infty} f(t)/g(t) \geq 1$ and $f(t) \lesssim g(t)$ if $\overline{\lim}_{t\to\infty} f(t)/g(t) \leq 1$. Furthermore, we say that $f(t) = o(g(t))$ if $\lim_{t\to\infty} f(t)/g(t) = 0$ and $f(t) = O(g(t))$ if $\overline{\lim}_{t\to\infty} f(t)/g(t) < \infty$. Also, we use the standard definition of an inverse function $f^{\leftarrow}(x) \triangleq \inf\{y : f(y) > x\}$ for a non-decreasing function $f(x)$; note that the notation $f(x)^{-1}$ represents $1/f(x)$. We use $\vee$ to denote $\max$, i.e., $x \vee y \equiv \max\{x, y\}$, and $\wedge$ to denote $\min$, i.e., $x \wedge y \equiv \min\{x, y\}$.

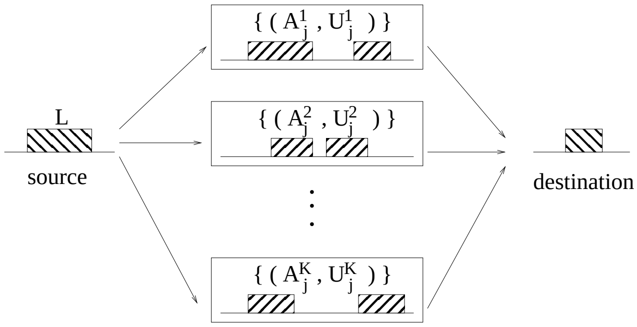

Fig. 1. Multipath transmission over K channels with failures

A. Single path transmission

For the case $K = 1$, there is only a single transmission path in the system, hence we let $A^1 \equiv A$. The total number of transmissions N and transmission time $T = T_r = T_s$ has been studied in [8], [9], [2].

Below we quote Propositions II.1 and II.2 from [8], [9], 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 II.1. *If there exists $\alpha > 0$ such that*

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

then,

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[N > n]}{\log n} = -\alpha. \tag{1}$$

Additionally, if $\mathbb{E}\left[U^{(\alpha\vee 1)+\theta}\right] < \infty$, $\mathbb{E}\left[A^{1+\theta}\right] < \infty$ *and* $\mathbb{E}\left[L^{\alpha+\theta}\right] < \infty$ *for some* $\theta > 0$*, then,*

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

Proposition II.2. *If*

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

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

$$\mathbb{P}[N > n] \sim \frac{\Gamma(\alpha + 1)}{\Phi(n)}, \tag{3}$$

and, under the same conditions as in Proposition II.1, as $t \to \infty$*,*

$$\mathbb{P}[T > t] \sim \frac{\Gamma(\alpha+1)(\mathbb{E}[U+A])^{\alpha}}{\Phi(t)}. \tag{4}$$

Remark. Proposition II.2 provides more refined results than Proposition II.1 under more restrictive conditions. One can easily check that (3) and (4) imply (1) and (2) by taking logarithms.

III. REDUNDANT TRANSMISSION

In this section we study redundant transmissions. We begin with K homogeneous paths, which is followed by the study of the general case of heterogenous paths. We investigate whether sending packets over K paths can mitigate the power law suffered from single path transmission.

A. Homogeneous paths

In this part, we present results for homogeneous paths. We first consider packets of the same size, and then study the more realistic case where packet sizes can be variable.

Proposition III.1. *If all packets are of constant size* $L \equiv l$ *and* $U \equiv 0$*, then,*

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

where γ *is the solution of* $\int_0^l e^{\gamma x} d\mathbb{P}[A \le x] = 1$.

This result can be easily derived using Corollary 3.2 in [2]. From this result, we see that using redundant transmission for equal size packets greatly improves performance, since the decay rate of the delay distribution increases as K increases, and thus in this case we obtain order improvements in delay performance when using redundant routing. In reality, however, packets are not of equal size. We next present a theorem for the case where the packet size is a random variable.

Theorem 1. If

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

$\mathbb{E}[L^{\alpha+\theta}] < \infty$, $\mathbb{E}\left[U^{(1\vee\alpha)+\theta}\right] < \infty$ and $\mathbb{E}\left[A^{1+\theta}\right] < \infty$ for some $\theta > 0$, then,

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

Remark. Comparing the above theorem and Proposition II.1, we observe that, the power law exponent of the total transmission time under redundant transmission is the same as that under single path transmission. Informally speaking, this is because $T_1, T_2, \ldots, T_K$ are not independent, since packets sent over these paths are of the same size.

This theorem is a direct consequence of Theorem 2, which investigates a more general scenario.

B. Heterogenous paths

For heterogenous paths, we have the following result when using redundant transmission.

Theorem 2. If

$$\lim_{x\to\infty} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \alpha_j \tag{5}$$

for $1 \le j \le K$, and $\alpha^* \triangleq \max_{1\le j\le K} \alpha_j > 0$, then, under the following three conditions I)-III), for some $\theta > 0$,

I) $\mathbb{E}[L^{\alpha+\theta}] < \infty$,

II) $\max_{1\le j\le K} \mathbb{E}\left[\left(U^j\right)^{(1\vee\alpha)+\theta}\right] < \infty$, and

III) $\max_{1\le j\le K} \mathbb{E}\left[\left(A^j\right)^{1+\theta}\right] < \infty$,

we have

$$\lim_{t\to\infty} \frac{\log \mathbb{P}[T_r > t]}{\log t} = -\alpha^*. \tag{6}$$

Remark. The above theorem implies 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).

of Theorem 2. First, we establish a lower bound by constructing a new system that has longer available periods than those found on all of the K paths. The construction is as follows. The new system has an on-off channel characterized by alternating i.i.d. sequences $\{\bar{A}_i\}$ and $\{\bar{U}_i\}$, where

$$\bar{A}_i = \max_{1\le j\le K} A_i^j$$

and $\bar{U}_i = 0$. Denote by $\underline{N}$ the number of transmissions of a packet of length L over this newly constructed channel.

Now, since $A_i^j, 1 \le j \le K$ are independent, we obtain

$$\mathbb{P}[\bar{A}_i > x] = 1 - \prod_{j=1}^{K} \mathbb{P}[A_i^j \le x].$$

Therefore,

$$\lim_{x\to\infty} \frac{\mathbb{P}[\bar{A}_i > x]}{\sum_{i=1}^{K} \bar{G}_j(x)} = 1,$$

coupled with (5), yields

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

which, by Proposition II.1, yields

$$\lim_{n\to\infty} \frac{\log \mathbb{P}[\underline{N} > n]}{\log n} = -\alpha^*. \tag{7}$$

Define $\underline{A}_i = \min_{1\le j\le K} A_i^j$ and $X_i \triangleq \underline{A}_i \mathbf{1}(x_1 < \underline{A}_i < x_2)$. Choosing x_1, x_2 such that $\mathbb{E}[X_i] > 0$, we obtain

$$T_r \ge \sum_{i=1}^{\underline{N}-1} X_i + L. \tag{8}$$

Therefore,

$$\begin{aligned}
\mathbb{P}\left[T_r > \frac{t}{\log t}\right] &\ge \mathbb{P}\left[\sum_{i=1}^{\underline{N}-1} X_i > \frac{t}{\log t}\right] \\
&\ge \mathbb{P}\left[\sum_{i=1}^{\underline{N}-1} X_i > \frac{t}{\log t}, \underline{N} > t\right] \\
&\ge \mathbb{P}\left[\underline{N} > t\right] - \mathbb{P}\left[\sum_{i=1}^{\underline{N}-1} X_i \le \frac{t}{\log t}, \underline{N} > t\right] \\
&\ge \mathbb{P}\left[\underline{N} > t\right] - \mathbb{P}\left[\sum_{i=1}^{\lfloor t\rfloor} X_i \le \frac{t}{\log t}\right].
\end{aligned} \tag{9}$$

Since $\mathbb{E}[e^{\theta X_i}] < \infty$ for some $\theta > 0$, we obtain, by a Chernoff bound, for some $\eta > 0$,

$$\begin{aligned}
&\mathbb{P}\left[\sum_{i=1}^{\lceil t\rceil} X_i \le t/\log t\right] \\
&\qquad \le \mathbb{P}\left[\sum_{i=1}^{\lceil t\rceil} \left(\mathbb{E}[X_i] - X_i\right) \ge \left(\mathbb{E}[X_i] - \frac{1}{\log t}\right) t\right] \\
&\qquad \le O\left(e^{-\eta t}\right),
\end{aligned} \tag{10}$$

which, in combination with (7) and (9), implies

$$\varliminf_{t\to\infty} \frac{\log \mathbb{P}[T_r > t]}{\log t} \ge -\alpha^*. \tag{11}$$

Next, we prove the upper bound. Since $\alpha^* \triangleq \max_{1\le j\le K} \alpha_i > 0$, there exists $1 \le j \le K$ such that $\alpha_j = \alpha^*$. For the jth path, we have $T_r \le T_j$ since $T_r = \min\{T_1, T_2, \cdots, T_K\}$. Using Proposition II.1, we obtain

$$\varlimsup_{t\to\infty} \frac{\log \mathbb{P}[T_r > t]}{\log t} \le \lim_{t\to\infty} \frac{\log \mathbb{P}[T_j > t]}{\log t} = -\alpha^*. \tag{12}$$

By combining (11) and (12), we complete the proof. □

Our preceding result characterizes the performance in terms of the "logarithmic asymptotics". Basically, it only contains information about the power law exponent, but yields no information about the pre-factor before the power law term. As a consequence, this result cannot distinguish between redundant transmission and single path transmission. In order to investigate the performance improvement for redundant transmission, we need a more refined asymptotic result. For a set of regularly varying functions $\Phi_j(\cdot)$, $1 \leq j \leq K$, we can compute the exact asymptotic tail of the distribution of T_r.

Theorem 3. If $\bar{F}(x)^{-1} \sim \Phi_j\left(\bar{G}_j(x)^{-1}\right)$ and

$$\lim_{x\to\infty} \frac{\Phi_j(x)}{\Phi(x)} = \zeta_j > 0, \tag{13}$$

where $\Phi(\cdot)$ is regularly varying with index $\alpha > 0$, then, under the conditions I)-III) in Theorem 2, as $t \to \infty$,

$$\mathbb{P}[T_r > t] \sim \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K} \left(\mathbb{E}[A^j + U^j]\right)^{-1} \zeta_j^{1/\alpha}\right)^{\alpha}} \frac{1}{\Phi(t)}. \tag{14}$$

Remark. From the preceding result, we see that, redundant transmission improves the system performance by reducing the tail of the distribution by a constant factor. If these K channels are i.i.d., this constant is equal to K^{α}.

In order to prove the theorem, we need the following lemma.

Lemma 1. For $\eta_j > 0$, $1 \leq j \leq K$,

$$\begin{aligned} &\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t] \\ &\sim \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K} \eta_j \zeta_j^{1/\alpha}\right)^{\alpha}} \frac{1}{\Phi(t)}. \end{aligned} \tag{15}$$

Proof. Observe that

$$\begin{aligned} &\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t] \\ &= \mathbb{E}\left[\prod_{j=1}^{K} \left(1 - \bar{G}_j(L)\right)^{\lfloor \eta_j t \rfloor}\right]. \end{aligned} \tag{16}$$

We can assume that $\Phi(x)$, $1 \leq j \leq K$ is absolutely continuous and strictly monotone since, by Proposition 1.5.8 of [3], one can always find an absolutely continuous and strictly monotone function, for c large enough,

$$\Phi^*(x) = \alpha \int_c^x \Phi(s) s^{-1} ds, \; x \geq c, \tag{17}$$

which satisfies

$$\bar{F}(x)^{-1} \sim \Phi_j\left((\bar{G}_j(x))^{-1}\right) \sim \zeta_j \Phi^*\left((\bar{G}_j(x))^{-1}\right).$$

Since $\Phi(x)$ is eventually non-decreasing, there exists x_0 such that for all $x > x_0$, $\Phi(x)$ has an inverse function $\Phi^{\leftarrow}(x)$. The condition (13) implies that, for $0 < \epsilon < 1$, there exists x_ϵ, such that for $x > x_\epsilon$ and $1 \leq j \leq K$,

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

and thus, by choosing $x_\epsilon > x_0$, we obtain

$$\Phi^{\leftarrow}\left((1-\epsilon)\bar{F}(x)^{-1}/\zeta_j\right) \leq (\bar{G}_j(x))^{-1}$$
$$\leq \Phi^{\leftarrow}\left((1+\epsilon)\bar{F}(x)^{-1}/\zeta_j\right). \tag{18}$$

First, we prove the lower bound. For $c > 0$ and x_ϵ selected in (18), choosing $x_t > x_\epsilon$ with $\Phi^{\leftarrow}\left(\bar{F}(x_t)\right) = t/c$, we obtain

$$\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t] \\
&\geq \mathbb{E}\left[\prod_{j=1}^{K}\left(1-\bar{G}_j(L)\right)^{\eta_j t}\mathbf{1}(L > x_t)\right] \\
&\geq \mathbb{E}\Bigg[\prod_{j=1}^{K}\left(1-\frac{1}{\Phi^{\leftarrow}\left((1-\epsilon)\bar{F}^(L)^{-1}/\zeta_j\right)}\right)^{\eta_j t} \\
&\qquad \cdot \mathbf{1}(\bar{F}(L) < \bar{F}(x_t))\Bigg],
\end{aligned}$$

which, by noting that $\bar{F}(L) \equiv V$ is a uniform random variable on $[0, 1]$, yields

$$\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t] \\
&\geq \mathbb{E}\left[\prod_{j=1}^{K}\left(1-\frac{1}{\Phi^{\leftarrow}\left((1-\epsilon)V^{-1}/\zeta_j\right)}\right)^{\eta_j t}\mathbf{1}(V < \bar{F}(x_t))\right].
\end{aligned} \tag{19}$$

Using Theorem 1.5.12 of [3], we know that $\Phi^{\leftarrow}(\cdot)$ is also regularly varying with index $1/\alpha$, which implies, for $\zeta > 0$,

$$\lim_{x\to\infty}\frac{\Phi^{\leftarrow}(\zeta x)}{\Phi^{\leftarrow}(x)} = \zeta^{1/\alpha},$$

and therefore, for x_t large enough and $V < \bar{F}(x_t)$,

$$\begin{aligned}
\frac{(1-\epsilon)^{1+1/\alpha}}{\zeta_j^{1/\alpha}}\Phi^{\leftarrow}\left(V^{-1}\right) &\leq \Phi^{\leftarrow}\left((1-\epsilon)V^{-1}/\zeta_j\right) \\
\frac{(1+\epsilon)^{1+1/\alpha}}{\zeta_j^{1/\alpha}}\Phi^{\leftarrow}\left(V^{-1}\right) &\geq \Phi^{\leftarrow}\left((1+\epsilon)V^{-1}/\zeta_j\right).
\end{aligned} \tag{20}$$

Letting $z = t/\Phi^{\leftarrow}\left(V^{-1}\right)$ and using (19), we obtain

$$
\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t] \\
&\geq \mathbb{E}\left[\prod_{j=1}^{K}\left(1 - \frac{\zeta_j^{1/\alpha}}{(1-\epsilon)^{1+1/\alpha}}\frac{1}{\Phi^{\leftarrow}\left(V^{-1}\right)}\right)^{\eta_j t}\mathbf{1}(V < \bar{F}(x_t))\right] \\
&\geq \int_0^c \prod_{j=1}^{K}\left(1 - \frac{\zeta_j^{1/\alpha}}{(1-\epsilon)^{1+1/\alpha}}\frac{z}{t}\right)^{\eta_j t}\frac{1}{\Phi\left(t/z\right)}\frac{\Phi'\left(t/z\right)}{\Phi\left(t/z\right)}\frac{t}{z^2}dz. \qquad (21)
\end{aligned}
$$

Since $\Phi(t)$ is regularly varying, by Theorem 1.4.1 and Theorem 1.2.1 of [3], it is easy to obtain uniformly for $0 \leq z \leq c$, as $t \to \infty$,

$$
\frac{\Phi(t)}{\Phi\left(t/z\right)} \sim z^{\alpha}
$$

and, recalling (17),

$$
\frac{\Phi'\left(t/z\right)}{\Phi\left(t/z\right)} = \frac{z\alpha}{t}.
$$

From the preceding results and (21), we obtain, as $t \to \infty$,

$$
\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t]\Phi(t) \\
&\sim \int_0^c \alpha e^{-z\left(\sum_{i=1}^{K}\eta_j\zeta_j^{1/\alpha}\right)(1-\epsilon)^{-1-1/\alpha}} z^{\alpha-1}dz \\
&= \frac{(1-\epsilon)^{\alpha+1}}{\left(\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}\right)^{\alpha}}\alpha\int_0^{c\left(\frac{\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}}{(1-\epsilon)^{1+1/\alpha}}\right)} e^{-x}x^{\alpha-1}dx,
\end{aligned}
$$

which, by passing $c \to \infty$ and $\epsilon \to 0$, yields

$$
\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t, N_2 > \eta_2 t, \cdots, N_K > \eta_K t]\Phi(t) \\
&\gtrsim \frac{1}{\left(\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}\right)^{\alpha}}\int_0^{\infty}\alpha e^{-x}x^{\alpha-1}dx \\
&= \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}\right)^{\alpha}}. \qquad (22)
\end{aligned}
$$

Next, we prove the upper bound. Observe that

$$
\begin{aligned}
&\mathbb{P}[N_1 > \eta_1 t + 1, N_2 > \eta_2 t + 1, \cdots, N_K > \eta_K t + 1] \\
&\leq \mathbb{E}\left[\prod_{j=1}^{K}\left(1 - \bar{G}_j(L)\right)^{\eta_j t}\mathbf{1}(L > x_t)\right] \\
&\quad + \mathbb{E}[\left(1 - \bar{G}_1(L)\right)^{\eta_1 t}\mathbf{1}(L < x_t)],
\end{aligned}
$$

which, by recalling (20) and using $1-x\le e^{x}$, yields

$$
\begin{aligned}
&\mathbb{P}[N_1>\eta_1 t+1, N_2>\eta_2 t+1,\cdots,N_K>\eta_K t+1]\\
&\le \mathbb{E}\left[\prod_{j=1}^{K}\left(1-\frac{\zeta_j^{1/\alpha}}{(1+\epsilon)^{1+1/\alpha}}\frac{1}{\Phi^{\leftarrow}\left(V^{-1}\right)}\right)^{\eta_j t}\mathbf{1}(V<\bar{F}(x_t))\right]\\
&\quad+o\left(\frac{1}{t^{\alpha+\epsilon}}\right)\\
&\le \mathbb{E}\left[e^{-\left(\frac{\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}}{(1+\epsilon)^{1+1/\alpha}}\right)\frac{t}{\Phi^{\leftarrow}\left(V^{-1}\right)}}\right]+o\left(\frac{1}{\Phi(t)}\right).
\end{aligned}
$$

For a constant integer $c>0$, $\zeta\triangleq\sum_{j=1}^{K}\eta_j\zeta_j^{1/\alpha}$ and $z=\zeta t/\left((1+\epsilon)^{1+1/\alpha}\Phi^{\leftarrow}\left(V^{-1}\right)\right)$, the preceding upper bound is no larger than

$$
\begin{aligned}
&\mathbb{E}\left[e^{-\frac{\zeta}{(1+\epsilon)^{1+1/\alpha}}\frac{t}{\Phi^{\leftarrow}\left(V^{-1}\right)}}\mathbf{1}\left(\frac{\zeta}{(1+\epsilon)^{1+1/\alpha}}\frac{t}{\Phi^{\leftarrow}\left(V^{-1}\right)}\le c\right)\right]\\
&\quad+\sum_{m=c}^{\infty}e^{-m}\mathbb{P}\left[m\le\frac{\zeta}{(1+\epsilon)^{1+1/\alpha}}\frac{t}{\Phi^{\leftarrow}\left(V^{-1}\right)}\le m+1\right]\\
&\quad+o\left(\frac{1}{\Phi(t)}\right)\\
&\le\int_0^c e^{-z}\left(\frac{\Phi'\left(\zeta t/((1+\epsilon)^{1+1/\alpha}z)\right)}{\Phi^2\left(\zeta t/\left((1+\epsilon)^{1+1/\alpha}z\right)\right)}\frac{\zeta t}{z^2(1+\epsilon)^{1+1/\alpha}}\right)dz\\
&\quad+\sum_{m=c}^{\infty}e^{-m}\frac{1}{\Phi\left(\frac{\zeta t}{(m+1)(1+\epsilon)^{1+1/\alpha}}\right)}+o\left(\frac{1}{\Phi(t)}\right).
\end{aligned}
$$

Since $\Phi(x)$ is regularly varying with index $\alpha>0$, we can choose c large enough, such that

$$
\frac{\Phi(t)}{\Phi\left(\frac{\zeta t}{(m+1)(1+\epsilon)^{1+1/\alpha}}\right)}\le(1+\epsilon)^{\alpha+2}\left(\frac{m+1}{\zeta}\right)^{\alpha}
$$

for all $m > c$. Therefore, as $t \to \infty$, we obtain

$$\mathbb{P}[N_1 > \eta_1 t + 1, N_2 > \eta_2 t + 1, \cdots, N_K > \eta_K t + 1]\Phi(t)$$
$$\leq \int_0^c e^{-z} \left(\frac{\Phi(t)\Phi'\left(\zeta t/((1+\epsilon)^{1+1/\alpha} z)\right)}{\Phi^2\left(\zeta t/\left((1+\epsilon)^{1+1/\alpha} z\right)\right)} \frac{\zeta t}{z^2(1+\epsilon)^{1+1/\alpha}} \right) dz$$
$$+ (1+\epsilon)^{\alpha+2} \sum_{m=c}^{\infty} e^{-m} \left(\frac{m+1}{\zeta} \right)^{\alpha} + o(1)$$
$$\lesssim \frac{(1+\epsilon)^{\alpha+1}}{\zeta^{\alpha}} \int_0^c \alpha e^{-z} z^{\alpha-1} dz$$
$$+ (1+\epsilon)^{\alpha+2} \sum_{m=c}^{\infty} e^{-m} \left(\frac{m+1}{\zeta} \right)^{\alpha},$$

which, by passing $\epsilon \to 0$ and $c \to \infty$, yields

$$\mathbb{P}[N_1 > \eta_1 t + 1, N_2 > \eta_2 t + 1, \cdots, N_K > \eta_K t + 1]\Phi(t)$$
$$\lesssim \frac{1}{\left(\sum_{j=1}^{K} \eta_j \zeta_j^{1/\alpha}\right)^{\alpha}} \int_0^{\infty} \alpha e^{-x} x^{\alpha-1} dx$$
$$= \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K} \eta_j \zeta_j^{1/\alpha}\right)^{\alpha}}. \tag{23}$$

Combining (20) and (23) finishes the proof. □

of Theorem 3. We begin with the upper bound. For $0 < \epsilon < 1$ and $\eta_j = 1/\mathbb{E}[A^j + U^j]$, we obtain,

$$\mathbb{P}[T_r > (1+2\epsilon)t] = \mathbb{P}\left[\bigcap_{j=1}^{K} \{T_j > (1+2\epsilon)t\} \right]$$
$$= \mathbb{P}\left[\bigcap_{j=1}^{K} \left\{ \sum_{i=1}^{N_j - 1} \left(A_i^j + U_i^j\right) + L > (1+2\epsilon)t \right\} \right]$$
$$\leq \mathbb{P}\left[\bigcap_{j=1}^{K} \left\{ \sum_{i=1}^{N_j} \left(A_i^j + \mathbb{E}[U^j]\right) > t \right\} \right]$$
$$+ \mathbb{P}\left[\bigcup_{j=1}^{K} \left\{ \sum_{i=1}^{N_j} \left(U_i^j - \mathbb{E}[U^j]\right) > \epsilon t \right\} \right]$$
$$+ \mathbb{P}[L > \epsilon t]. \tag{24}$$

Then, using union bound, we derive

$$\begin{aligned}
&\mathbb{P}[T_r > (1+2\epsilon)t] \leq \\
&\quad \mathbb{P}\left[\bigcap_{j=1}^{K}\left\{\sum_{i=1}^{N_j}\left(A_i^j + \mathbb{E}[U^j]\right) > t, N_j > (1-\epsilon)\frac{t}{\mathbb{E}[A^j+U^j]}\right\}\right] \\
&\quad + \sum_{j=1}^{K}\mathbb{P}\left[\sum_{i=1}^{N_j}\left(A_i^j + \mathbb{E}[U^j]\right) > t, N_j \leq (1-\epsilon)\frac{t}{\mathbb{E}[A^j+U^j]}\right] \\
&\quad + \sum_{j=1}^{K}\mathbb{P}\left[\left\{\sum_{i=1}^{N_j}\left(U_i^j - \mathbb{E}[U^j]\right) > \epsilon t\right\}\right] \\
&\quad + \mathbb{P}[L > \epsilon t] \\
&\quad \leq \mathbb{P}\left[\bigcap_{j=1}^{K}\{N_j > (1-\epsilon)\eta_j t\}\right] \\
&\qquad + \sum_{j=1}^{K}\mathbb{P}\left[\sum_{i=1}^{(1-\epsilon)\eta_j t}\left(A_i^j \wedge L + \mathbb{E}[U^j]\right) > t\right] \\
&\qquad + \sum_{j=1}^{K}\mathbb{P}\left[\left\{\sum_{i=1}^{N_j}\left(U_i^j - \mathbb{E}[U^j]\right) > \epsilon t\right\}\right] \\
&\qquad + \mathbb{P}[L > \epsilon t] \\
&\quad \triangleq I_1 + I_2 + I_3 + I_4.
\end{aligned} \tag{25}$$

Using the result (4.20) in [9], we know $I_2 + I_3 + I_4 = o\left(1/\Phi(t)\right)$, which, in view of Lemma 1, yields

$$\mathbb{P}[T_r > t] \lesssim \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K}\left(\mathbb{E}[A^j+U^j]\right)^{-1}\zeta_j^{1/\alpha}\right)^{\alpha}}\frac{1}{\Phi(t)}. \tag{26}$$

Next, we proceed with proving the lower bound.

$$\mathbb{P}[T_r > t] \geq \mathbb{P}\left[\bigcap_{j=1}^{K}\left\{\sum_{i=1}^{N_j-1}\left(A_i^j + U_i^j\right) > t\right\}\right]$$

$$
\begin{aligned}
&\geq \mathbb{P}\left[\bigcap_{j=1}^{K}\left\{N_j > (1-\epsilon)\frac{t}{\mathbb{E}[A^j+U^j]}\right\}\right] \\
&-\sum_{j=1}^{K}\mathbb{P}\left[\sum_{i=1}^{N_j-1}\left(A_i^j+U_i^j\right) > t, N_j \leq (1-\epsilon)\frac{t}{\mathbb{E}[A^j+U^j]}\right] \\
&\triangleq I_1 - I_2. \qquad (27)
\end{aligned}
$$

Using the same approach in deriving (4.30) in [9], we can prove that $I_2 = o\left(1/\Phi(t)\right)$. Therefore, by Lemma 1, we obtain

$$
\mathbb{P}[T_r > t] \gtrsim \frac{\Gamma(\alpha+1)}{\left(\sum_{j=1}^{K}\left(\mathbb{E}[A^j+U^j]\right)^{-1}\zeta_j^{1/\alpha}\right)^{\alpha}}\frac{1}{\Phi(t)},
$$

which, in combination with (26), finishes the proof. □

IV. Split transmission

Next, we study the case when a packet is split into several pieces and sent over K independent paths. Since each path has to successfully transmit a fragment of the original packet, we are interested in the total number of retransmissions $N = \sum_{j=1}^{K} N_j$ and the transmission time $T_s \triangleq \max_{1\leq j\leq K} T_j$. Using the derived results, we will determine which of the two strategies, split transmission or redundant transmission, results in a lighter distribution tail.

We begin with homogeneous paths, and then investigate heterogenous paths. A fraction γ_j of the packet L is sent over path j, $\sum_{j=1}^{K}\gamma_j = 1$, $0 \leq \gamma_j \leq 1, 1 \leq j \leq K$. We derive the optimal splitting strategy that minimizes the exponent of the transmission time tail.

A. Homogeneous paths

We have the following theorem for split transmission over homogenous paths, where each packet is evenly split into K pieces. Its proof is a special case of that for heterogeneous paths (see Theorem 5), and hence is omitted.

Theorem 4. Under the same conditions in Theorem 1, if there exists $\beta > 0$, such that

$$
\lim_{x\to\infty}\frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \beta, \qquad (28)
$$

then,

$$
\lim_{t\to\infty}\frac{\log P(T_s > t)}{\log t} = -\beta\alpha.
$$

Remark. Since $\beta \geq 1$, comparing the results in Proposition II.1 and Theorem 1, we see that, for homogeneous paths, split transmission is no worse than redundant transmission when packets are split evenly.

To provide some concrete examples, we next consider several typical distributions, and compute the power law exponent for these distributions.

1) Examples for typical distributions: Theorem 4 indicates that the effectiveness of split transmission is closely related to the packet size distribution, as characterized by (28). We next apply this result to several families of distributions to further illustrate this point. For each distribution, we calculate α and β, and the power law tail exponent is equal to $-\beta\alpha$.

- Exponential distribution. Consider the case where the size of the packet, L, follows an exponential distribution with parameter λ and the available time period, A, follows an exponential distribution with parameter μ, i.e.,

$$\begin{aligned}\bar{F}(x) &= P(L > x) = e^{-\lambda x},\\ \bar{G}(x) &= P(A > x) = e^{-\mu x}.\end{aligned}$$

Then,

$$\begin{aligned}\alpha &= \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \frac{\log\left(e^{-\lambda x}\right)}{\log\left(e^{-\mu x}\right)} = \frac{\lambda}{\mu},\\ \beta &= \frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \frac{\log\left(e^{-\lambda K x}\right)}{\log\left(e^{-\lambda x}\right)} = \frac{\lambda K}{\lambda} = K.\end{aligned}$$

- Weibull distribution. Consider

$$\begin{aligned}\bar{F}(x) &= P(L > x) = e^{-(\lambda x)^b},\\ \bar{G}(x) &= P(A > x) = e^{-(\mu x)^b},\end{aligned}$$

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

$$\begin{aligned}\alpha &= \frac{\log \bar{F}(x)}{\log \bar{G}(x)} = \frac{\log e^{-(\lambda x)^b}}{\log e^{-(\mu x)^b}} = \frac{-(\lambda x)^b}{-(\mu x)^b} = \left(\frac{\lambda}{\mu}\right)^b,\\ \beta &= \frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \frac{\log\left(e^{-(\lambda K x)^b}\right)}{\log\left(e^{-(\lambda x)^b}\right)} = \frac{(\lambda K)^b}{(\lambda)^b} = K^b.\end{aligned}$$

- Pareto distribution. Consider

$$\bar{F}(x) = \begin{cases} (b_1/x)^\lambda, & x \geq b_1\\ 1, & x < b_1,\end{cases}$$

$$\bar{G}(x) = \begin{cases} (b_2/x)^\mu, & x \geq b_2\\ 1, & x < b_2\end{cases}$$

where $\lambda > 0, \mu > 0$, and $b_1, b_2 > 0$. Then,

$$\begin{aligned}\alpha &= \lim_{x\to\infty} \frac{\lambda(\log b_1 - \log x)}{\mu(\log b_2 - \log x)} = \frac{\lambda}{\mu},\\ \beta &= \lim_{x\to\infty} \frac{\lambda(\log b_1 - \log K - \log x)}{\lambda(\log b_1 - \log x)} = 1.\end{aligned}$$

Remark. Observe that $\beta = 1$ when both F and G follow Pareto distributions. In that case, split transmission has the same power law exponent as the single path transmission and redundant transmission. Furthermore, Pareto distribution is not the only type of distribution under which split transmission is not beneficial. We illustrate this point in the next subsection.

2) When is split transmission not beneficial?: We next show another family of distributions under which split transmission is not beneficial. Consider the following distribution

$$\bar{F}(x;\eta,\xi) = \begin{cases} 1, & x \in (-\infty, 1] \\ e^{-\eta(\log x)^{\xi}}, & x \in [1, +\infty) \end{cases}$$

where $\eta > 0$, $\xi > 0$. Then

$$\frac{\log \bar{F}(Kx)}{\log \bar{F}(x)} = \left(1 + \frac{\log K}{\log x}\right)^{\xi} \to 1, \quad \text{as } x \to \infty.$$

Compute the p-th moment of $\bar{F}(x)$ ($p > 0$),

$$\begin{aligned} m_p(\eta,\xi) &= -\int x^p d\bar{F}(x;\eta,\xi) \\ &= -\int_1^{\infty} x^p de^{-\eta(\log x)^{\xi}} \\ &= -\int_0^{\infty} e^{py} de^{-\eta y^{\xi}} \\ &= \eta\xi \int_0^{\infty} y^{\xi-1} e^{py-\eta y^{\xi}} dy. \end{aligned}$$

There are three regions for parameter ξ,

- $\xi = 1$. $\bar{F}(x;\eta,\xi)$ is a Pareto distribution and m_p exists iff $p < \eta$.
- $\xi < 1$. $\bar{F}(x;\eta,\xi)$ has a heavier tail than any power law, and m_p does not exists for any p.
- $\xi > 1$. $\bar{F}(x;\eta,\xi)$ has a lighter tail than any power law, and m_p exists for all $p > 0$.

Note that for large ξ, the probability tail can decay fast, albeit slower than exponential. In this case, we still do not have a gain using split transmission in terms of increasing the power law exponent.

The above discussion indicates that the Pareto distribution is not the only type of distribution under which split transmission leads to no benefits in mitigating power law delays. In fact, distributions that do not benefit from split transmission can have either heavier or lighter tails than Pareto distributions. Split transmission is not beneficial as long as $\beta = 1$ in (28), e.g., when $\log\left(1/\bar{F}(x)\right)$ is slowly varying.

B. Heterogenous paths

For heterogenous paths, a packet of size L is split into K smaller fragments of sizes $\gamma_1 L, \gamma_2 L, \ldots, \gamma_K L$, respectively, where $\sum_{j=1}^{K} \gamma_j = 1$, $0 \le \gamma_j \le 1, 1 \le j \le K$. We have the following result on packet transmission delay.

Theorem 5. If there exist α_j, β_j, $j = 1, 2, \ldots, K$ such that

$$\lim_{x\to\infty} \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \alpha_j, \tag{29}$$

$$\lim_{x\to\infty} \frac{\log \bar{F}(x)}{\log \bar{F}(\gamma_j x)} = \beta_j, \tag{30}$$

with $\alpha^\circ \triangleq \min_{1\le j\le K} \beta_j\alpha_j > 0$, then,

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

and, under the conditions I)-III) in Theorem 2,

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

Remark. When paths are heterogeneous, the packet transmission delay is determined by the best paths under redundant transmission and by the worst paths under split transmission. On the other hand, split transmission only sends a fraction of the packet on each path. Comparing this to Theorem 2, we see that, if $\min_{1\le j\le K} \beta_j\alpha_j > \max_{1\le j\le K} \alpha_j$, split transmission is more beneficial than redundant transmission in minimizing the tail behavior; otherwise, redundant transmission is more beneficial. We will show later that, by carefully choosing the way to split packets, split transmission can always result in tail performance that is no worse than redundant transmission.

of Theorem 5. We begin with proving the result for T_s. Since

$$T_s = \max_{1\le j\le K} T_j,$$

we obtain, using a union bound,

$$\max_{1\le j\le K} \mathbb{P}[T_j > t] \le \mathbb{P}\left[T_s > t\right] \le \sum_{j=1}^{K} \mathbb{P}[T_j > t], \tag{31}$$

Next, using (29) and (30), we derive

$$\lim_{x\to\infty} \frac{\log P(\gamma_j L > x)}{\log P(A_j > x)} = \lim_{x\to\infty} \frac{\beta_k \log \bar{F}(x)}{\log \bar{G}_k(x)} = \beta_j\alpha_j,$$

which, by Proposition II.1, yields

$$\lim_{t\to\infty} \frac{\log P(T_j > t)}{\log t} = -\beta_j\alpha_j.$$

Thus, for $\epsilon > 0$, there exists $t_0 > 0$ such that for all $t > t_0$,

$$-\beta_j\alpha_j - \epsilon < \frac{\log \mathbb{P}[T_j > t]}{\log t} < -\beta_j\alpha_j + \epsilon.$$

Hence, for $t > t_0$, we have

$$\max_{1\le j\le K} \mathbb{P}[T_j > t] > t^{-\alpha^\circ - \epsilon}$$

and

$$\sum_{j=1}^{K} \mathbb{P}[T_j > t] < K t^{-\alpha^\circ + \epsilon},$$

which, combined with (31) and passing $\epsilon \to 0$, yields

$$\lim_{t\to\infty} \frac{\log P(T_s > t)}{\log t} = -\min_{1\le j\le K} \{\beta_j \alpha_j\} = -\alpha^\circ.$$

Now, we derive the result for N. Since

$$N_s = \sum_{j=1}^{K} N_j,$$

we have

$$\max_{1\le j\le K} \mathbb{P}[N_j > n] \le \mathbb{P}[N > n] \le \sum_{j=1}^{K} \mathbb{P}\left[N_j > \frac{n}{K}\right]. \tag{32}$$

Proposition II.1 implies

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

which, combined with (32) and using a similar argument as in proving the result for T_s, yields

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

□

1) Optimal split transmission: From Theorem 5, we can see that in order to optimize the power law delay tail, we need to choose $\gamma_1, \gamma_2, \ldots, \gamma_K$ so that $\min_{1\le j\le K} \beta_j \alpha_j$ is maximized. To achieve this, we may speculate that we need to choose $\gamma_1, \gamma_2, \ldots, \gamma_K$ so that $\beta_1\alpha_1 = \beta_2\alpha_2 = \cdots = \beta_K\alpha_K$. The following theorem confirms that this is true when $\log\left(1/\bar{F}(x)\right)$ is not slowly varying.

Theorem 6. Suppose we use split transmission over K heterogeneous paths, each satisfying (29). 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 (i) there exists a unique constant $\rho \geq 0$ such $\beta(\gamma) = \gamma^{-\rho}$; and (ii) the optimal splitting scheme that minimizes the power law exponent of $\mathbb{P}[T_s > t]$ satisfies:

a) If $\rho > 0$, then

$$\gamma_j^* = \frac{\alpha_j^{1/\rho}}{\sum_{i=1}^K \alpha_i^{1/\rho}}. \tag{33}$$

b) If $\rho = 0$, then let $\gamma_j = 0$ for $\alpha_j \neq \max_{1\leq j\leq K} \alpha_j$ and the other γ_j can take arbitrary values.

The corresponding optimal power law exponent for $\mathbb{P}[T_s > t]$ is $-\alpha_\rho$, where

$$\alpha_\rho = \begin{cases} \left(\sum_{i=1}^K \alpha_i^{1/\rho}\right)^\rho, & \rho > 0, \\ \max_{1\leq j\leq K} \alpha_j, & \rho = 0. \end{cases} \tag{34}$$

Remark. In the preceding result, we only minimize the power law exponent. When $\rho = 0$, we have $\beta(\gamma) = 1$, and $\log\left(1/\bar{F}(x)\right)$ is a slowly varying function. In this case, we should only use the best paths, and the scheme in (33) is to split arbitrarily among the best paths. For this case, we need a more refined asymptotic result that accounts for not only the power law exponent but also the exact pre-factors to derive the optimal split strategy. Due to limited space, we do not study this problem. When $\rho > 0$, all the channels are utilized, and the optimal fraction on each path is specified by (33). In this case, one can easily check that the optimal tail exponent is indeed achieved when $\beta_1\alpha_1 = \beta_2\alpha_2 = \cdots = \beta_K\alpha_K$.

Remark. Note that $\alpha_\rho = \left(\sum_{i=1}^K \alpha_i^{1/\rho}\right)^\rho \geq \alpha^*$ with equality if and only if $\rho = 0$, where $\alpha^* = \max_{1\leq j\leq K} \alpha_j > 0$, as defined in Theorem 2. Thus, under the assumption of Theorem 5, split transmission achieves a better exponent than redundant transmission if $\rho > 0$.

of Theorem 6. (i) 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 $\beta_0 = \beta(\gamma_0) > 1$ 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 integer m, n,

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

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

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

or, equivalently,

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

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

(ii) Let $\{\gamma_j^*\}$ be an optimal split scheme and $-\alpha_\rho$ the corresponding optimal exponent. By Theorem 5,

$$\alpha_\rho = \min_{j:\gamma_j>0} \alpha_j(\gamma_j^*)^{-\rho}. \tag{35}$$

If $\rho = 0$, then

$$\alpha_\rho = \min_{j:\gamma_j>0} \alpha_j \le \max_{1\le j\le K} \alpha_j = \alpha^*$$

with equality if and only if $\gamma_j = 0$ whenever $\alpha_j \neq \alpha^*$.

If $\rho > 0$, then (35) gives

$$\gamma_j^*(\alpha_\rho)^{1/\rho} \le \alpha_j^{1/\rho}, \quad j = 1, 2, \ldots, K.$$

Summing over j and noting $\sum_j \gamma_j^* = 1$, we have

$$(\alpha^*)^{1/\rho} \le \sum_{j=1}^{K} \alpha_j^{1/\rho}$$

with equality if γ_j^* is given by (33). □

2) Optimal split transmission examples: To illustrate the results obtained in the preceding section, we compute the optimal split transmission scheme for some typical distributions.

- Exponential distribution. Consider the case where the size of the packet, L, follows an exponential distribution with parameter λ and the available time period on path j, A^j, follows an exponential distribution with parameter μ_j, i.e.,

  $$\bar{F}(x) = P(L > x) = e^{-\lambda x},$$

  $$\bar{G}_j(x) = P(A^j > x) = e^{-\mu_j x}.$$

  Then,

  $$\alpha_j = \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \frac{\lambda}{\mu_j}$$

  $$\beta(\gamma) = \frac{\log \bar{F}(x)}{\log \bar{F}(\gamma x)} = \frac{1}{\gamma},$$

  and

  $$\rho = -\log \beta(\gamma)/\log \gamma = 1.$$

  Therefore, the optimal split is

  $$\gamma_j = \frac{\frac{\lambda}{\mu_j}}{\sum_{i=1}^{K} \frac{\lambda}{\mu_i}} = \frac{\frac{1}{\mu_j}}{\sum_{i=1}^{K} \frac{1}{\mu_i}}, \quad j = 1, \ldots, K.$$

- Weibull distribution. Consider the case where the size of the packet, L, and the available time period on path j,

A^j, follow Weibull distributions, i.e.,

$$\bar{F}(x) = P(L > x) = e^{-(\lambda x)^b},$$

$$\bar{G}_j(x) = P(A^j > x) = e^{-(\mu_j x)^b},$$

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

$$\alpha_j = \frac{\log \bar{F}(x)}{\log \bar{G}_j(x)} = \frac{-(\lambda x)^b}{-(\mu_j x)^b} = \left(\frac{\lambda}{\mu_j}\right)^b,$$

$$\beta(\gamma) = \frac{\log \bar{F}(x)}{\log \bar{F}(\gamma x)} = \frac{1}{\gamma^b},$$

and

$$\rho = -\log \beta(\gamma)/\log \gamma = b.$$

Therefore, the optimal split is

$$\gamma_j = \frac{\frac{\lambda}{\mu_j}}{\sum_{i=1}^K \frac{\lambda}{\mu_i}} = \frac{\frac{1}{\mu_j}}{\sum_{i=1}^K \frac{1}{\mu_i}}, \quad j = 1, \ldots, K.$$

- Pareto distribution. Consider the case where the size of the packet, L, and the available time period on path j, A^j, follow Pareto distributions. In this case, we have $\beta(\gamma) = 1$. The optimal split transmission strategy is to split among the best paths.

3) Exact asymptotic result for optimal split transmission: Our proposed optimal split transmission minimizes the power law exponent of $\mathbb{P}[T_s > t]$. In other words, Theorem 6 only characterizes the tail behavior in the logarithmic scale. Next, to refine the result, we present a theorem on the exact asymptotic result for optimal split transmission.

Theorem 7. If $\log(\bar{F}(x)^{-1}) = x^\rho l(x)$ where $\rho > 0$ and $l(x)$ is slowly varying with

$$e^{l(x)} \sim e^{l(\gamma x)} \tag{36}$$

for $\gamma > 0$, and

$$\bar{F}(x)^{-1} \sim \zeta_j \left(\Phi\left(\bar{G}_j(x)^{-1}\right)\right)^{\alpha_j/\alpha}, \tag{37}$$

where $\alpha_j, \zeta_j > 0$ and $\Phi(\cdot)$ is regularly varying with index $\alpha > 0$, then, under the conditions I)-III) in Theorem 2, as $t \to \infty$,

$$\mathbb{P}[T_s > t] \sim$$

$$\sum_{l=1}^K (-1)^{l+1} \sum_{\{j_1,\cdots,j_l\} \subseteq \{1,\cdots,K\}} \frac{\Gamma(\alpha_\rho + 1)}{\left(\sum_{s=1}^l \eta_{j_s} \zeta_{j_s}^{\frac{1}{\alpha_{j_s}}}\right)^{\alpha_\rho}} \frac{1}{\Phi(t)^{\frac{\alpha_\rho}{\alpha}}}, \tag{38}$$

where $\eta_j \triangleq 1/\mathbb{E}[A^j + U^j]$, $\alpha_\rho \triangleq \left(\sum_{j=1}^K \alpha_j^{1/\rho}\right)^\rho$.

In order to prove the theorem, we need the following lemma.

Lemma 2. For $\eta_j > 0$, $1 \leq j \leq K$,

$$1 - \mathbb{P}[N_1 < \eta_1 t, N_2 < \eta_2 t, \cdots, N_K < \eta_K t]$$

$$\sim \sum_{l=1}^K (-1)^{l+1} \sum_{\{j_1,\cdots,j_l\}\subseteq\{1,\cdots,K\}} \frac{\Gamma(\alpha_\rho + 1)}{\left(\sum_{s=1}^l \eta_{j_s} \zeta_{j_s}^{\frac{1}{\alpha_{j_s}}}\right)^{\alpha_\rho}} \frac{1}{\Phi(t)^{\frac{\alpha_\rho}{\alpha}}}. \tag{39}$$

Proof. Observe that

$$\begin{aligned} &1 - \mathbb{P}[N_1 < \eta_1 t, N_2 < \eta_2 t, \cdots, N_K < \eta_K t] \\ &= 1 - \mathbb{E}\left[\prod_{j=1}^K \left(1 - \left(1 - \bar{G}_j(\gamma_j L)\right)^{\lfloor \eta_j t \rfloor}\right)\right] \\ &= \sum_{l=1}^K (-1)^{l+1} \sum_{1\leq j_1<\cdots<j_l\leq K} \mathbb{E}\left[\prod_{s=1}^l \left(1 - \bar{G}_{j_s}(\gamma_{j_s} L)\right)^{\lfloor \eta_{j_s} t \rfloor}\right], \end{aligned} \tag{40}$$

and therefore, we only need to compute, for $l \geq 1$ and $\{j_1, \cdots, j_l\} \subseteq \{1, \cdots, K\}$,

$$\mathbb{E}\left[\prod_{s=1}^l \left(1 - \bar{G}_{j_s}(\gamma_{j_s} L)\right)^{\lfloor \eta_{j_s} t \rfloor}\right].$$

Similar to the proof of Lemma 1, we can assume that $\Phi(x)$, $1 \leq j \leq K$ is absolutely continuous and strictly monotone. The conditions (36) and (37) imply that, for $0 < \epsilon < 1, \gamma > 0$, there exists x_ϵ, such that for $x > x_\epsilon$ and $1 \leq j \leq K$,

$$(1-\epsilon)\left(\bar{F}(x)\right)^{-\gamma^\rho} \leq \zeta_j \left(\Phi(\bar{G}_j(\gamma x)^{-1})\right)^{\frac{\alpha_j}{\alpha}} \leq (1+\epsilon)\left(\bar{F}(x)\right)^{-\gamma^\rho},$$

and thus, by choosing $x_\epsilon > x_0$, we obtain

$$\begin{aligned} \Phi^{\leftarrow}\left(\left(\frac{(1-\epsilon)}{\zeta_j}\right)^{\alpha/\alpha_j} \bar{F}(x)^{-\alpha\gamma^\rho/\alpha_j}\right) &\leq \bar{G}_j(\gamma x)^{-1} \\ &\leq \Phi^{\leftarrow}\left(\left(\frac{(1+\epsilon)}{\zeta_j}\right)^{\alpha/\alpha_j} \bar{F}(x)^{-\alpha\gamma^\rho/\alpha_j}\right). \end{aligned} \tag{41}$$

Using (41) and the same approach in deriving (19), we obtain,

$$
\begin{aligned}
&\mathbb{E}\left[\prod_{s=1}^{l}\left(1-\bar{G}_{j_s}(\gamma_{j_s}L)\right)^{\lfloor \eta_{j_s} t\rfloor}\right] \\
&\geq \mathbb{E}\left[\prod_{s=1}^{l}\left(1-\bar{G}_{j_s}(\gamma_{j_s}L)\right)^{\eta_{j_s} t}\mathbf{1}(L>x_t)\right] \\
&\geq \mathbb{E}\Bigg[\prod_{s=1}^{l}\left(1-\frac{1}{\Phi^{\leftarrow}\left(((1-\epsilon)/\zeta_{j_s})^{\alpha/\alpha_{j_s}}\,V^{-\alpha\gamma_{j_s}^{\rho}/\alpha_{j_s}}\right)}\right)^{\eta_{j_s} t} \\
&\qquad \cdot \mathbf{1}(V<\bar{F}(x_t))\Bigg],
\end{aligned}
$$

which, by defining $\eta(x) \triangleq \Phi(x)^{\alpha_\rho/\alpha}$ and noting that $\gamma_{j_s}^{\rho}/\alpha_{j_s} = 1/\alpha_\rho$, yields

$$
\begin{aligned}
&\mathbb{E}\left[\prod_{s=1}^{l}\left(1-\bar{G}_{j_s}(\gamma_{j_s}L)\right)^{\lfloor \eta_{j_s} t\rfloor}\right] \\
&\geq \mathbb{E}\Bigg[\prod_{s=1}^{l}\left(1-\frac{1}{\eta^{\leftarrow}\left(((1-\epsilon)/\zeta_{j_s})^{\alpha_\rho/\alpha_{j_s}}\,V^{-1}\right)}\right)^{\eta_{j_s} t} \\
&\qquad \cdot \mathbf{1}(V<\bar{F}(x_t))\Bigg].
\end{aligned}
$$

Following the same procedure in computing (21), we obtain

$$
\begin{aligned}
&\mathbb{E}\left[\prod_{s=1}^{l}\left(1-\bar{G}_{j_s}(\gamma_{j_s}L)\right)^{\lfloor \eta_{j_s} t\rfloor}\right] \\
&\qquad\qquad \gtrsim \frac{\Gamma(\alpha_\rho+1)}{\left(\sum_{s=1}^{l}\eta_{j_s}\zeta_{j_s}^{1/\alpha_{j_s}}\right)^{\alpha_\rho}}\frac{1}{\Phi(t)^{\alpha_\rho/\alpha}}.
\end{aligned}
\tag{42}
$$

Repeating a similar procedure in deriving the upper bound for Lemma 1, we can prove

$$
\begin{aligned}
&\mathbb{E}\left[\prod_{s=1}^{l}\left(1-\bar{G}_{j_s}(\gamma_{j_s}L)\right)^{\lfloor \eta_{j_s} t\rfloor}\right] \\
&\qquad\qquad \lesssim \frac{\Gamma(\alpha_\rho+1)}{\left(\sum_{s=1}^{l}\eta_{j_s}\zeta_{j_s}^{1/\alpha_{j_s}}\right)^{\alpha_\rho}}\frac{1}{\Phi(t)^{\alpha_\rho/\alpha}},
\end{aligned}
\tag{43}
$$

which, in combination with (42), finishes the proof of the lemma. □

of Theorem 7. We begin with the upper bound. For $0 < \epsilon < 1$ and $\eta_j = 1/\mathbb{E}[A^j + U^j]$, we obtain,

$$\begin{aligned}\mathbb{P}[T_s > (1+2\epsilon)t] &= \mathbb{P}\left[\bigcup_{j=1}^{K} \{T_j > (1+2\epsilon)t\}\right] \\ &\leq \mathbb{P}\left[\bigcup_{j=1}^{K}\left\{\sum_{i=1}^{N_j}\left(A_i^j + \mathbb{E}[U^j]\right) > t\right\}\right] \\ &\quad + \mathbb{P}\left[\bigcup_{j=1}^{K}\left\{\sum_{i=1}^{N_j}\left(U_i^j - \mathbb{E}[U^j]\right) > \epsilon t\right\}\right] \\ &\quad + \mathbb{P}[L > \epsilon t]. \end{aligned} \tag{44}$$

Then, using union bound, we derive

$$\begin{aligned}&\mathbb{P}\left[\bigcup_{j=1}^{K}\left\{\sum_{i=1}^{N_j}\left(A_i^j + \mathbb{E}[U^j]\right) > t\right\}\right] \\ &\leq \mathbb{P}\left[\bigcup_{j=1}^{K} \{N_j > (1-\epsilon)\eta_j t\}\right] \\ &\quad + \sum_{j=1}^{K}\mathbb{P}\left[\sum_{i=1}^{(1-\epsilon)\eta_j t}\left(A_i^j \wedge L + \mathbb{E}[U^j]\right) > t\right]. \end{aligned} \tag{45}$$

Combining (44), (45) and using the results in deriving (26), we obtain, by Lemma 2,

$$\begin{aligned}&\mathbb{P}[T_s > t] \lesssim \mathbb{P}\left[\bigcup_{j=1}^{K} \{N_j > \eta_j t\}\right] \\ &\sim \sum_{l=1}^{K}(-1)^{l+1} \sum_{\{j_1,\cdots,j_l\}\subseteq\{1,\cdots,K\}} \frac{\Gamma(\alpha_\rho + 1)}{\left(\sum_{s=1}^{l} \eta_{j_s}\zeta_{j_s}^{\frac{1}{\alpha_{j_s}}}\right)^{\alpha_\rho}} \frac{1}{\Phi(t)^{\frac{\alpha_\rho}{\alpha}}}. \end{aligned}$$

Using a similar approach as in deriving the lower bound for Theorem 3, we can derive a lower bound that coincides with the upper bound for this theorem, which completes the proof. □

V. Conclusion

Parallelism is a common approach to improve reliability and efficiency in practice. In this paper, we investigate whether and how parallelism can be used to improve network performance. Specifically, we study whether and how multipath transmission can mitigate power law delays. We show that, when all packets are of the same size, redundant transmission can greatly reduce the transmission delay in the sense that the ratio of the delay distribution tail with and without redundant transmission tends to zero. However, when packet sizes are random variables such that $\log \mathbb{P}[L > x] \approx \alpha^* \log \mathbb{P}[A^j > x]$, we prove that, maybe counter intuitively, redundant transmission cannot change the order of the probability tail of the transmission delays, and can only improve the system performance by a constant factor. We also show that split transmission is effective in mitigating power 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. Last, we provide an optimal split transmission strategy when the paths are heterogeneous, and further derive an exact asymptotic result for packet delivery time under this scheme. Our results can be extended to many other applications 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.

REFERENCES

[1] L. N. Andersen and S. Asmussen. Parallel comuputing, failure recovery and extreme values. *Journal of Statistical Theory and Applications*, 2008. Forthcoming.

[2] S. Asmussen, P. Fiorini, L. Lipsky, T. Rolski, and R. Sheahan. Asymptotic behavior of total times for jobs that must start over if a failure occurs. *Mathematics of Operations Research*, 33(4):932–944, November 2008.

[3] N. H. Bingham, C. M. Goldie, and J. L. Teugels. *Regular Variation*, volume 27. Cambridge University Press, 1987.

[4] D. G. Cantor and M. Gerla. Optimal transmission in a packet-switched computer network. *IEEE Transactions on Computers*, 23(10), October 1974.

[5] R. Gallager. A minimum delay transmission algorithm using distributed computation. *IEEE Transactions on Communications*, 25(1), 1977.

[6] P. R. Jelenković and J. Tan. Is ALOHA causing power law delays? In *Proceedings of the 20th International Teletraffic Congress*, Ottawa, Canada, June 2007*; Lecture Notes in Computer Science,* No 4516, pp. 1149-1160, Springer-Verlag, 2007.

[7] P. R. Jelenković and J. Tan. Are end-to-end acknowlegements causing power law delays in large multi-hop networks? In *14th Informs Applied Probability Conference*, Eindhoven, July 9-11 2007.

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

[9] P. R. Jelenković and J. Tan. Characterizing heavy-tailed distributions induced by retransmissions. Technical report, Department of Electrical Engineering, Columbia University, EE2007-09-07, September 2007. Eprint arXiv: 0709.1138v2.

[10] L. Kleinrock. *Communication Nets: Stochastic Message Flow and Delay*. New York: McGraw-Hill Book Company, 1964.

[11] T. S. Rappaport. *Wireless communications: principles and practise*. Prentice Hall, 2 edition, January 2002.

[12] 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.