Estimating message transmission time over heterogeneous disrupted links

Philip Ginzboorg
Huawei Technologies, and
Aalto University, Finland
Email: philip.ginzboorg@iki.fi

Valtteri Niemi
University of Helsinki, Finland, and
Helsinki Institute for Information Technology
Email: valtteri.niemi@helsinki.fi

Jörg Ott
Technical University of Munich, Germany and
Aalto University, Finland
Email: ott@in.tum.de

***Abstract*—We consider fragmented message transmission through a heterogeneous chain of several independently disrupted communication links. The message is prepared for fragmentation before transmission by dividing it into blocks of constant size. In this setting, we derive an approximation for the mean and standard deviation of fragmented message transmission time when one of the links in the heterogeneous chain is much slower than the rest.**

***Index Terms*—data transfer, fragmentation, channel with failures, DTN.**

I. Introduction

The system we consider in this paper is illustrated in Figure 1. It is a chain of n disruptive communication links that connect $n+1$ nodes. The first and last nodes in the chain are the sender and receiver nodes, respectively. The intermediate nodes forward (relay) the data traffic from sender to receiver. This chain is initially empty, except for a message waiting for transmission opportunity in the first (sender) node, where, before transmission, the message is prepared for fragmentation by dividing it into blocks of constant size f.

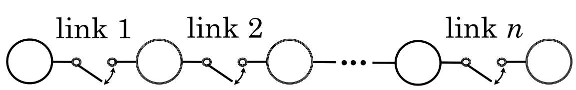

Fig. 1. Schematic illustration of a chain of n disrupted links.

A link changes its state between connected (ON) and disconnected (OFF) in a random manner and independently from the states in the other links; the ON and OFF state durations of a link may have different distributions; and it is assumed that the distributions of disruptions in each link do not change with time. We shall denote by X the duration of ON state, and by Y the duration of OFF state in a link. The message size x and the block size f also have the units of time. This is justified, because we assume that the transmission speed during ON state is constant and the same for all links. We refer the reader to [1], where transmission time of fragmented message of a given size x over a single disrupted link was studied by the authors, for a further discussion of this link model.

The papers [2]–[5] complement [1], in that they consider how the distribution of message size x may affect the distribution of transmission time over a single disrupted link. In previous work [6], we studied the transmission time of fragmented messages over a chain of homogeneous disrupted links, where the ON and OFF, respectively, state durations of different links follow the same distribution.

In this paper, we relax the homogeneity assumption and consider a chain with one "bottleneck" link; the OFF state durations for that link tend to be (much) longer than in other links. We assume that n is not large: Path lengths are often (i) observed to be fairly short in disruption-tolerant networks (DTN) and opportunistic networks of workable scale [7]; or (ii) are designed to be short for performance reasons in multihop wireless networks [8], [9]. Our main contribution are new estimates of mean and standard deviation (or variance) of transmission time of fragmented messages over such chain.

II. Estimating mean transmission time

The idea of the mean approximation is simple. Let us first consider the case where the bottleneck is the last link in the chain. Typically, there is a queue in the node before the bottleneck. For the approximation, we assume that after the time point when the first block has arrived to the bottleneck there is indeed all the time a queue before the bottleneck.

Now we proceed as follows. We estimate the transmission time of the first block over each non-bottleneck link and sum up these values. Then we add an estimate for the time it takes to transmit the whole message over the bottleneck link.

Let us next consider the case where the bottleneck is the first link in the chain. When the last block has finally passed the bottleneck, the rest of the chain is typically fairly empty. For the approximation, we assume that it is completely empty. We estimate the time it takes to transmit the whole message over the bottleneck link, and add the sum of estimates of transmission times of the last block over all non-bottleneck links. It can be noticed that the final result is the same as in the previous case, although the justification logic is different.

Finally, the case where the bottleneck is in the middle of the chain would lead to the same approximation, when both justification logics are applied together.

Let us denote by T_o the time it would take on the average for one *single* block to transfer through the chain *assuming* the bottleneck link is removed from the chain. The second part of the mean approximation is an estimate for the time it takes for the *whole* message to transfer through the bottleneck link *assuming* all other links are removed from the chain. This time is denoted by T_b.

Now the approximation T for mean transmission time of the whole message over the whole chain can be computed as the sum of the two partial estimates:

$$T = T_o + T_b. \tag{1}$$

Let us consider a chain where all non-bottleneck links have identically distributed ON-OFF patterns; the only link that has a different distribution is the bottleneck link. In this case, the estimate T_o is obtained by simply multiplying the corresponding estimate for a single (non-bottleneck) link with the number of these links; (this number is of course one less than the total number of links in the chain because the bottleneck link was removed); thus $T_o = (n-1)T_1$, and

$$T = (n-1)T_1 + T_b. \tag{2}$$

where T_1 is the expected time it takes one block to transfer over one non-bottleneck link.

In the rest of this section we show how T_1 and T_b can be estimated if we know the distributions of ON and OFF state durations in each link. Since both T_1 and T_b are estimates in the setting of a single disrupted link, we can use approximations developed in [1].

Let us begin with T_1, and illustrate the approach with a concrete example, where uniformly distributed ON and OFF state durations take values in the interval $[0\,\mathrm{s}, 2\,\mathrm{s}]$, with a mean of $1\,\mathrm{s}$.

One block with length f passes one link as follows: Its arrival time is either during OFF state duration or ON state duration, each with probability 1/2. The expected length of the state durations can be computed: In both cases it is $4/3\,\mathrm{s}$. (This is an instance of the classical "waiting time paradox" [10], [11].)

In case the arrival is during OFF state, the expected remainder of this OFF state duration is $2/3\,\mathrm{s}$. Let us denote the expected remaining transfer time after the OFF state ends with u. The following ON state lasts long enough to accommodate the block with probability $P(X \geq f) = (2\,\mathrm{s} - f)/(2\,\mathrm{s})$. In that case the additional time from this ON state that is needed to transfer the block through the link is f.

With $f/(2\,\mathrm{s})$ probability the ON state duration is too short to accommodate the block. The expected length of too short ON state duration is $f/2$. Next we have to wait for another OFF state duration, the expected length of which is $1\,\mathrm{s}$. After the OFF state duration is over, the remaining time is still u.

Now we can conclude that

$$u = \frac{2\,\mathrm{s} - f}{2\,\mathrm{s}} f + \frac{f}{2\,\mathrm{s}} \left(\frac{f}{2} + 1\,\mathrm{s} + u \right),$$

and we can solve this equation for u.

Hence, the expected transfer time is $2/3\,\mathrm{s} + u$ under the condition that the block arrives when the link is OFF.

Let us next consider that the block arrives when the link is ON. The expected duration of this ON state is $4/3\,\mathrm{s}$. Now we can have a reasonable estimate that the remainder of the ON state duration is uniformly distributed between $0\,\mathrm{s}$ and $4/3\,\mathrm{s}$. Hence, the remainder is long enough to accommodate the block of length f with approximate probability $(4/3\,\mathrm{s} - f)/(4/3\,\mathrm{s})$. If this is the case, the transfer time is simply f.

With approximate probability $f/(4/3\,\mathrm{s})$ the remainder is too short. Then the required time is $f/2 + 1\,\mathrm{s} + u$.

Putting all together, the closed formula for T_1 is

$$T_1 = \left[-9f^2 + 70f + 16\right] / \left[24(2-f)\right]. \tag{3}$$

where the time unit is second (s).

As another example, let us estimate T_1 when lengths of the ON and OFF state durations are exponentially distributed with mean of $1\,\mathrm{s}$. The calculation in this case is rather simple because the remainder of the first state duration has mean length of $1\,\mathrm{s}$. (This is because the exponential distribution is memoryless.)

When the block arrives, the link is ON with probability $1/2$. If the block arrives during OFF state, it has to wait $1\,\mathrm{s}$ on the average (the remaining part of OFF state duration), after which we get to the same situation as when the block arrives during an ON state. If we denote the expected transfer time when the arrival is during ON state with u, then the corresponding mean transfer time when the arrival is during OFF state is $1\,\mathrm{s} + u$. Hence, we can conclude that the expected transfer time is $(1/2)u + (1/2)(1\,\mathrm{s} + u) = u + 1/2\,\mathrm{s}$.

Next we compute u given f. A block of length f fits into an ON state duration with probability $P(X \geq f) = e^{-f/(1\,\mathrm{s})}$. For simplicity, we will omit below the division by $1\,\mathrm{s}$ in the exponential function. Hence, we can conclude that

$$u = e^{-f} \cdot f + (1 - e^{-f})(w + 1\,\mathrm{s} + u)$$

where w is the expected length of a too short ON epoch. Indeed, after the too short ON epoch we have to wait for one OFF epoch after which we are back in "square one" and the remaining time is still u. Because we have exponential distribution $w < f/2$. On the other hand, especially for small values of f, the difference $f/2 - w$ is very small. For simplicity, we use the estimate $w = f/2$ in the derivation below.

Now we can solve $u \approx f/2 - 1\,\mathrm{s} + e^f(f/2 + 1\,\mathrm{s})$. The total mean transfer time over one non-bottleneck link can now be approximated by

$$T_1 = \left(f + e^f(f + 2\,\mathrm{s}) - 1\,\mathrm{s}\right)/2. \tag{4}$$

It is straight-forward to generalize formulas (3) and (4) to uniform and exponential distributions with other parameter values. We have also derived a formula for T_1 in case of a general distribution. The formula and its derivation are in the Appendix.

Let us turn to T_b, the estimate of the expected transfer time of the whole message over the bottleneck link.

The calculation of T_b is based on Wald's identity [12, p. 103]:

$$T_b = E(N) \cdot (E(X) + E(Y)), \tag{5}$$

where $E(N)$ is the mean number of ON-OFF cycles that is needed to transmit message of size x over the bottleneck link, and $E(X) + E(Y)$ is the average duration of one ON-OFF cycle.

We estimate $E(N)$ as follows. Let us define X_f as the "useful" part of ON state duration, i.e. the biggest multiple of the block size f that can fit into X: $X_f = \lfloor X/f \rfloor \cdot f$, where $\lfloor a \rfloor$ is the integer part of a. After we have sent as many blocks as we could fit into X, there may still remain some ON state time that is smaller than f. Let W_f be that "wasted" part of ON state duration:

$$W_f = X - X_f. \tag{6}$$

Our estimate of $E(N)$ is

$$E(N) = x/E(X_f) = x/(E(X) - E(W_f)), \tag{7}$$

where x is the message size.

Let us illustrate the approach with a concrete example, where we consider uniformly distributed X and Y, with $E(X) = 1\,\text{s}$, and $E(Y) = 2\,\text{s}$. The block size is $f = 0.1\,\text{s}$.

In this case, W_f is uniformly distributed in the interval $[0, 0.1)$ and thus $E(W_f) = 0.05\,\text{s}$. Therefore, by (7), $E(N) = 8\,\text{s}/(1\,\text{s} - 0.05\,\text{s}) = 8.42$; and by (5), $T_b = 8.42 \cdot 3\,\text{s} = 25.3\,\text{s}$.

For $E(Y) = 4\,\text{s}$ we get $T_b = 8.42 \cdot 5\,\text{s} = 42.1\,\text{s}$; for $E(Y) = 8\,\text{s}$ we get $T_b = 75.8\,\text{s}$; and for $E(Y) = 16\,\text{s}$ we get $T_b = 143\,\text{s}$.

Calculation of $E(W_f)$ in the case of exponentially distributed X are slightly different. In the example case of $f = 1\,\text{s}$, the $E(W_f)$ is the same as w above (because exponential distribution is memoryless). It can be computed that $E(X_f) = 1\,\text{s} - 0.418\,\text{s} = 0.582\,\text{s}$.

Different values of f can be handled with a similar reasoning, both in the case of uniform distribution and in the case of exponential distribution.

III. Estimating variance of transmission time

The basic idea of the variance approximation is also very simple. We ignore all the non-bottleneck links and try to estimate the variance of the transmission time of the whole message over the bottleneck link. This estimate is denoted by V_b and the variance approximation V for the whole message over the whole chain is:

$$V = V_b. \tag{8}$$

We justify using this simple approximation for variance as follows.

If there would be only one block in the message then the transmission times over different links would be independent and, therefore, the variance of the total transmission time t would be the sum of the variances over individual links. When the message contains several blocks the situation is different. For simplicity, let us discuss a case where the bottleneck link is the first one in the chain. Now t could be written as $t_b + t_l$ where t_b is the time it takes to transfer the whole message through the bottleneck and t_l is the additional time it takes to transfer the last block through the rest of the chain.

The random variables t_b and t_l are negatively correlated due to the following reason. If t_b is high then it is also likely that, when the last block finally transfers over the first (i.e. the bottleneck) link, the rest of the chain is relatively empty. Indeed, all the earlier blocks have had plenty of time to get through the chain of non-bottleneck links. In that case, the last block has to wait less in the queues before the non-bottleneck links for the earlier blocks to be transmitted, and it moves faster through the rest of the chain. On the other hand, if t_b is low, then it is likely that the last block will have to wait in the queues before the non-bottleneck links, and its transmission time over the rest of the chain becomes longer.

All in all, the variance of the total transmission time is less than the sum of variances of t_b and t_l.

Looking at different links, we can make an educated guess that the variance of transmission time over the bottleneck link is bigger than the variance of transmission time over a non-bottleneck link. Therefore, it is reasonable to try first an approximation where we simply ignore the non-bottleneck links and estimate the total variance with variance over the bottleneck.

The calculation of V_b is based on Blackwell-Girshick formula [12, p. 107]: if $S(X)$ is the sum of N i.i.d. random variables X, where N is also a random variable, then

$$\text{Var}(S(X)) = \text{Var}(N)[E(X)]^2 + E(N)\text{Var}(X). \tag{9}$$

In our case $S(X)$ is the sum of N ON state durations needed to transmit a message of size x over the bottleneck link. We know $E(X)$ and $\text{Var}(X)$ in that link, and mean number of ON states $E(N)$ can be calculated using methods outlined in the previous section.

What remains is $\text{Var}(N)$. This can be estimated by the following formula, mentioned in section IV of [1], that relates $\text{Var}(N)$ to the message size x and the first two moments of X_f, in case $x \gg E(X_f)$:

$$\text{Var}(N) \approx \text{Var}(X_f)x/[E(X_f)]^3. \tag{10}$$

Since $X_f = X - W_f$ by (6), we have that $E(X_f) = E(X) - E(W_f)$, and on the other hand, $X = X_f + W_f$ and

$$\text{Var}(X) = \text{Var}(X_f) + \text{Var}(W_f) + 2\text{Cov}(X_f, W_f).$$

We are able to approximate $\text{Var}(X_f)$ by neglecting $\text{Cov}(X_f, W_f)$, which amounts to assuming that X_f and W_f are independent:

$$\text{Var}(X_f) \approx \text{Var}(X) - \text{Var}(W_f). \tag{11}$$

All in all, if we know $E(W_f)$ and $\text{Var}(W_f)$, we are able to estimate $E(N)$ and $\text{Var}(N)$, and then $\text{Var}(S(X))$ by (9).

Please note that if f is small compared to a typical duration of ON state $E(X)$, then, independently of the distribution of X, it is reasonable to assume that the distribution of W_f resembles uniform in the interval $[0, f)$. Thus, in case of small f,

$$E(W_f) \approx f/2, \quad \text{and} \quad \text{Var}(W_f) \approx f^2/12. \tag{12}$$

The variance of total duration of OFF states $S(Y)$ is calculated using

$$\text{Var}(S(Y)) = \text{Var}(N)[E(Y)]^2 + E(N)\text{Var}(Y). \tag{13}$$

Because the durations of ON and OFF states are independent, we get the estimate V_b as the sum

$$V_b = \text{Var}(S(X)) + \text{Var}(S(Y)). \tag{14}$$

TABLE I

% ERROR FOR EXP. AND UNI. DISTRIBUTED LINK DISRUPTIONS; FIRST COLUMN IS THE MEAN OFF STATE DURATION IN THE BOTTLENECK LINK.

		exponentially distributed disruptions								uniformly distributed disruptions							
nr. of links n:		1		3		5		8		1		3		5		8	
$E(Y)$/(s)	f/(s)	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2	ϵ_1	ϵ_2
1	0.01	-2.9	-2.2	-20.1	-13.1	-38.0	-19.1	-45.2	-24.9	-1.0	-2.7	-28.4	-4.7	-28.5	-8.3	-35.5	-12.9
	0.1	-2.8	-1.9	-28.2	-14.7	-37.3	-20.3	-43.9	-25.9	-0.8	-2.4	-19.9	-7.4	-27.9	-11.2	-34.5	-16.0
	1	-1.7	1.5	-25.8	-29.4	-31.4	-35.8	-33.9	-42.1	2.7	0.2	-16.5	-30.4	-22.4	-36.3	-26.0	-42.2
2	0.01	-4.9	-2.6	-8.9	5.7	-25.5	9.3	-32.4	11.3	-1.8	-2.9	-18.0	5.1	-14.2	10.8	-19.9	16.1
	0.1	-4.9	-2.2	-18.0	5.0	-25.1	8.5	-31.6	10.7	-1.5	-2.6	-8.7	4.1	-13.7	9.3	-19.2	14.0
	1	-3.1	-0.6	-16.4	-7.5	-21.6	-7.5	-25.1	-10.0	2.6	0.1	-4.3	-11.4	-8.8	-11.2	-12.9	-12.8
4	0.01	-7.2	-2.8	-5.4	1.9	-17.0	4.8	-21.4	8.7	-2.5	-2.8	-12.7	-0.7	-7.8	1.3	-10.5	3.8
	0.1	-6.8	-2.4	-12.7	1.8	-16.5	4.9	-20.7	8.9	-2.1	-2.7	-5.1	-0.9	-7.2	1.0	-9.8	3.3
	1	-4.4	-1.0	-10.1	-2.5	-13.2	0.1	-16.0	3.0	2.5	0.1	0.8	-7.2	-0.7	-6.4	-2.6	-5.4
8	0.01	-8.9	-3.0	-4.6	-0.7	-13.3	0.0	-19.5	1.8	-3.3	-2.9	-11.2	-2.0	-5.7	-1.4	-6.9	-0.6
	0.1	-8.5	-2.7	-11.0	-1.1	-12.8	0.2	-14.9	2.0	-2.8	-2.6	-4.2	-2.1	-5.1	-1.4	-6.3	-0.7
	1	-5.3	-0.9	-7.6	-2.2	-9.0	-0.8	-10.2	1.0	2.3	0.3	2.0	-4.0	1.6	-3.8	1.0	-3.8
16	0.01	-9.7	-2.8	-4.3	-1.9	-11.9	-1.8	-12.9	-1.1	-3.6	-3.1	-10.9	-2.5	-4.9	-2.2	-5.6	-2.0
	0.1	-9.4	-2.5	-10.7	-2.0	-11.5	-1.5	-12.4	-0.8	-3.1	-2.5	-3.9	-2.6	-4.3	-2.2	-4.9	-2.0
	1	-6.0	-1.2	-6.9	-1.9	-7.4	-1.3	-7.9	-0.6	2.2	0.3	2.1	-2.0	2.1	-1.9	1.9	-2.1

IV. Simulation results

We have built a tailor-made program in the C programming language to simulate fragmented message transmission over a chain of disrupted links. The setting of parameters in those simulations included one value from each of the following sets:

- Number of links $n \in \{1, 3, 5, 8\}$;
- Distributions of state durations in the n links: $n \in$ {exponential, uniform}.
- mean(ON): $E(X) = 1\,\mathrm{s}$;
- mean(OFF): $E(Y) = 1\,\mathrm{s}$, if non-bottleneck;
- mean(OFF): $E(Y) \in \{2\,\mathrm{s}, 4\,\mathrm{s}, 8\,\mathrm{s}, 16\,\mathrm{s}\}$, if bottleneck;
- message size $x = 8\,\mathrm{s}$;
- block size $f \in \{0.01\,\mathrm{s}, 0.1\,\mathrm{s}, 1\,\mathrm{s}\}$.

We ran simulations for each combination of parameters: a sequence of 10^4 messages was transmitted through the chain; this was repeated with 20 different seeds, resulting in a set of $2 \cdot 10^5$ recorded transmission times t_i per each parameter permutation. The transmission of any message in a sequence is preceded by a random delay, during which the links have a chance to change state several times. The delay is exponentially distributed, with a mean that is 10 times the average duration of the ON-OFF cycle in the chain.

When post-processing the data, we have computed (a) the sample mean $\bar{t}$ and standard deviation s from the $2 \cdot 10^5$ data points t_i; (b) the estimates T and V, as described above, using equations from (3) to (14); (c) the averages α and β of $\bar{t}$ and s, respectively, for all different positions of the bottleneck link in the chain; and (d) the relative errors ϵ_1 and ϵ_2:

$$\epsilon_1 = (T - \alpha)/\alpha, \qquad \epsilon_2 = (\sqrt{V} - \beta)/\beta. \tag{15}$$

Errors in estimating mean and variance are calculated by comparing the estimates to the averages of $\bar{t}$ and s in (15), because our simple approximations do not take into account the position of the bottleneck link in the chain. In the cases we have simulated, the values of $\bar{t}$ and s are typically slightly different when the bottleneck is in the middle of the chain, compared to when it is at the beginning or end of the chain, but these differences are significantly smaller than the differences between those values and the estimates.

Table I shows ϵ_1 and ϵ_2 as percentage points; the shaded cells contain error values of more than 10%. We will now list some of the most important observations from the table.

(i) When $n = 1$ (bottleneck only), our approximations to mean and standard deviation are pretty good. Actually, if this would not be the case, it is unlikely that the approximations would work for longer chains either.

(ii) When $n > 1$ and there is no bottleneck (i.e. we have a homogeneous chain), our approximations do not work very well. Of course, this makes a lot of sense because the logic behind our approximations is based on the existence of a bottleneck. (We have included the homogeneous case to check what happens to the approximations when the difference between the bottleneck and the other links decreases.)

(iii) Related to the previous observation, we notice that, indeed, the errors get smaller as the severity of the bottleneck increases.

(iv) The approximations in case of uniformly distributed disruptions are, in general, better than in case of exponentially distributed disruptions.

(v) Rather surprisingly, our approximation for standard deviation has typically smaller errors than the approximation for mean.

V. Conclusion

We have developed simple approximations for the mean and standard deviation of transmission time of a fragmented message over a chain of disrupted links that contains a single bottleneck link. These approximations have been verified in cases where the times between disruptions are either uniformly or exponentially distributed. One main observation is that as the bottleneck link gets slower, the approximations get better.

For the cases of uniform and exponential distributions of disruptions, we have developed formulas to estimate mean and variance of message transmission time as a function of the message and fragmentation block sizes. Similar formulas could be derived and evaluated for other distributions, following the same logic. This holds for both theoretical distributions and for distributions obtained by measurements in real systems. Another extension is to verify how well the approximations work for very short or very long messages.

The approximations do not require that disruptions in non-bottleneck links have same statistics. In this case we would use a slightly modified version of equation (2).

Another way to extend our results is to consider situations where distributions of disruptions in different links are not known, but we can obtain (e.g., via measurements) approximate values for the mean and standard deviation of transmission time of fragmented message in each link. The mean and standard deviation of message transmission time over the whole chain could be then estimated based on (1) and (modified version of) (2). Real-life experiments in challenged networks, e.g., DTN, would be needed to verify how good this estimate is.

REFERENCES

[1] P. Ginzboorg, V. Niemi, and J. Ott, "Message fragmentation for disrupted links," in *Proceedings of European Wireless 2011*. Vienna, Austria: VDE VERLAG GMBH, April 27-29 2011, pp. 415–424.

[2] P. R. Jelenković and J. Tan, "Dynamic packet fragmentation for wireless channels with failures," in *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, ser. MobiHoc '08. New York, NY, USA: ACM, 2008, pp. 73–82. [Online]. Available: http://doi.acm.org/10.1145/1374618.1374629

[3] J. Nair, M. Andreasson, L. L. Andrew, S. H. Low, and J. C. Doyle, "File fragmentation over an unreliable channel," in *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–9.

[4] J. Tan, Swapna, and Shroff, "Retransmission delays with bounded packets: Power-law body and exponential tail," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 27–38, Feb. 2014. [Online]. Available: http://dx.doi.org/10.1109/TNET.2013.2244907

[5] J. Nair, M. Andreasson, L. L. H. Andrew, S. H. Low, and J. C. Doyle, "On channel failures, file fragmentation policies, and heavy-tailed completion times," *IEEE/ACM Trans. Netw.*, vol. 24, no. 1, pp. 529–541, Feb. 2016. [Online]. Available: http://dx.doi.org/10.1109/TNET.2014.2375920

[6] P. Ginzboorg, V. Niemi, and J. Ott, "Message fragmentation for a chain of disrupted links," *Computer Communications*, vol. 48, pp. 84–97, July 2014. [Online]. Available: http://dx.doi.org/10.1016/j.comcom.2014.03.015

[7] A. Chaintreau, A. Mtibaa, L. Massoulie, and C. Diot, "The Diameter of Opportunistic Mobile Networks," in *Proc. of ACM CoNEXT*, 2007.

[8] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla, "The Impact of Multihop Wireless Channel on TCP Throughput and Loss," in *Proc. of IEEE INFOCOM*, 2003.

[9] M. Petrova, L. Wu, M. Wellens, and P. Mähnen, "Hop of No Return: Practical Limitations of Wireless Multi-Hop Networking," in *Proceedings of RealMAN 2005*, July 2005.

[10] W. Feller, *An Introduction to Probability Theory and its Applications*, 2nd ed. New York: John Wiley & Sons, 1971, vol. II, ch. I.4.

[11] L. Kleinrock, *Queueing Systems, Volume 1: Theory*. Hoboken, NJ, USA: John Wiley & Sons, 1975, ch. 5.2.

[12] A. Klenke, *Probability Theory: a Comprehensive Course*. London, UK: Springer Verlag, 2008, ch. 5.1.

APPENDIX: CALCULATION OF T_1 IN GENERAL CASE

In this Appendix we derive a formula for computing T_1, the expected time it takes one block of size f to transfer over one non-bottleneck link, assuming that we know the distributions of X and Y in that link.

The block arrives at a link during ON part of the link's ON-OFF cycle with probability $p = E(X)/(E(X) + E(Y))$, and during OFF part of the cycle with probability $q = 1 - p$.

Let r be the time between the block arrival and until the next change of link state; if the arrival happens in the ON state, this "remaining" time will be denoted by r_X; and if the arrival happens in the OFF state, it will be denoted by r_Y.

The following formulas can be derived based on, e.g., [11, pp. 172-173, equations (5.10), (5.16)]:

$$P(r_X \geq f) = 1 - \int_0^f P(X > y)dy/E(X). \qquad (16)$$

$$E(r_Y) = \left(E(Y)^2 + \mathrm{Var}(Y)\right)/(2E(Y)). \qquad (17)$$

Let us denote by u the mean transmission time of the block if the counting of time starts at a beginning of ON state; it can be found by first step analysis: With probability $P(X \geq f)$, $u = f$; and with probability $P(X < f)$, $u = E(X \mid X < f) + E(Y) + u$. Therefore,

$$u = P(X \geq f)f + P(X < f)\left(E(X \mid X < f) + E(Y) + u\right),$$

and solving for u results in

$$u = f + \frac{P(X < f)}{P(X \geq f)} \cdot \left(E(X \mid X < f) + E(Y)\right). \qquad (18)$$

Please note that $E(X \mid X < f) + E(Y)$ is the average duration of ON-OFF cycle with a too short ON state, and it can be shown that $P(X < f)/P(X \geq f)$ is the mean number of such cycles before successful transmission occurs.

The term $E(X \mid X < f)$ in (18) can be computed by

$$E(X \mid X < f) = \int_0^f y p_X(y)dy/P(X < f), \qquad (19)$$

where p_X is the probability density function (PDF) of X.

Let us denote by u_X the mean transmission time of a block when it arrives during ON state. Similarly, let us denote by u_Y the mean transmission time of a block when it arrives during OFF state.

If the arrival happens during OFF state, then the block has to wait for $E(r_Y)$ time on the average, before its transmission starts from the beginning of ON state. Thus,

$$u_Y = E(r_Y) + u. \qquad (20)$$

If the arrival happens during ON state, then, with probability $P(r_X \geq f)$, the block transmission time is f, and with probability $P(r_X < f)$, there is a waiting time of $E(r_X \mid r_x < f) + E(Y)$ on the average, before transmission starts from the beginning of ON state. Thus,

$$\begin{aligned} u_X = {} & P(r_X \geq f) \cdot f + \qquad (21) \\ & P(r_X < f)\left(E(r_X \mid r_x < f) + E(Y) + u\right). \end{aligned}$$

The term $E(r_X \mid r_X < f)$ can be computed by

$$E(r_X \mid r_X < f) = \int_0^f y p_{r_X}(y)dy/P(r_X < f), \qquad (22)$$

where p_{r_X} is the PDF of r_X. This function can be computed, similarly as was done for (16), by $p_{r_X}(y) = P(X > y)/E(X)$.

Finally,

$$T_1 = q \cdot u_Y + p \cdot u_X. \qquad (23)$$