Computer Communications 48 (2014) 84–97

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Message fragmentation for a chain of disrupted links ☆

Philip Ginzboorg [a,b,*], Valtteri Niemi [c], Jörg Ott [b]

[a] Huawei Technologies OY, Itämerenkatu 9, FIN-00180 Helsinki, Finland
[b] Aalto University, Otakaari 5A, FIN-02150 Espoo, Finland
[c] University of Turku, Department of Mathematics and Statistics, FIN-20014 Turku, Finland

ARTICLE INFO

Article history:
Available online 27 March 2014

Keywords:
Fragmentation
Channel with failures
DTN

ABSTRACT

We investigate the problem of estimating the transmission time of fragmented messages over multiple disrupted links. We build a system model for the case where a single message is sent over a chain of links and the disruptions in these links are identically and independently distributed. For this case, we derive approximation formulas for the mean transmission time, based on number of links, length of fragments and distributions of disruptions. The formulas are verified against simulation experiments in the cases of uniform and exponential distributions for disruptions.

1. Introduction

A challenged network is a network subject to difficult operational constraints, like disrupted links and high delays. Potential applications of challenged networks are in the areas were infrastructure needed for good end-to-end connectivity is difficult or inconvenient to deploy. Those areas include communications in industrial environments (mines, factories, shipyards), space, military, emerging markets, and local opportunistic communication between mobile devices.

Ways of communicating in challenged environments are developed by the research community. One of the most encompassing and well documented efforts is the work done in the IETF Delay Tolerant Networking Research Group (DTNRG).[1] A delay-tolerant network (DTN) can be defined as a network that does not require for its operation (i) small Round-Trip Time (RTT), or (ii) simultaneous end-to-end paths, or (iii) continuous connectivity between nodes [1]. Since communication in challenged environments can be implemented from DTNRG specifications, we will use DTN notation in this paper. Please note, however, that our results are not restricted only to those challenged networks that follow the DTNRG specifications.

In a challenged network with unstable transmission links the connection between the sender and the receiver may be cut before the entire message has been transmitted. For that reason the *contact* times (i.e. the times when the link between two nodes is available, or: in the ON state) can be a very scarce resource. Allowing messages to be fragmented on their way to the destination may help to use these contact times better.

Delay-tolerant networks may use potentially large messages (rather than small packets) as basic transmission unit offered to applications. Here again, sending large messages implies that those will be broken down into individual packets for the actual transmission across a physical link that comply with the link's Maximum Transfer Unit (MTU) size. Such a mechanism is defined, e.g., for the convergence layers of the DTN bundle protocol [2].

Since messages may be large, their transmission as a series of packets may not complete during a contact period. When a link comes up again after a down ("OFF") period (the *inter-contact time*), the message transmission should resume (roughly) where it stopped, rather than have to restart from the beginning. For this purpose, it is required to *fragment* a message into smaller pieces ("units") whose transmission is more likely to fit into a contact period than the complete message.

Two types of fragmentation for DTN are defined in [2]: pro-active and reactive. In the former, the source node for the link divides application data into blocks and sends each block in a separate fragment. In the latter, the data is split only when the transmission between two nodes on any link of the message path is interrupted; resulting in one fragment with data that made it to the receiver and one containing the remainder at the sender. The fragmented data is re-assembled at its destination, but also an intermediate node can re-assemble fragments into a new, bigger piece.

We shall assume that the sender does message *quantization*, i.e. it prepares the message for fragmentation by dividing it into blocks

☆ Expanded version of a talk given at the AOC 2012 workshop in San Francisco, California, USA, on June 25 2012.

* Corresponding author at: Huawei Technologies OY, Itämerenkatu 9, FIN-00180 Helsinki, Finland. Tel.: +358 504836224.

E-mail addresses: philip.ginzboorg@iki.fi (P. Ginzboorg), valtteri.niemi@utu.fi (V. Niemi), joerg.ott@tkk.com (J. Ott).

[1] URL: http://www.dtnrg.org.

http://dx.doi.org/10.1016/j.comcom.2014.03.015
0140-3664/

of size f, called "fragmentation unit", and this is done before transmission. The message may be fragmented on its way to the destination only along the borders defined by f. This is motivated by security and efficiency considerations. Firstly, since a contact between two nodes may abruptly end, the sender (be it the originating or an intermediate node) must decide on the fragmentation borders and add the appropriate message authentication codes (MACs) *before* transmitting the message.

Secondly, marshaling and assembly of the message pieces at the destination is easier with fixed f [3].

Message quantization leads to the question "How does the transmission time of a message depend on f?" The simplest way to answer this question in an actual network is by trial and error. For example, in IP networks the maximum size of an IP packet that can be transmitted without fragmentation is typically determined by the path probing technique of RFC 1981 [4]. But it is hard to apply this technique in a challenged network where simultaneous end-to-end path from source to destination is unlikely. Another, complimentary way to answer this question is to estimate the dependency from the known network conditions. This is the kind of an answer that we investigate in this paper.

Recent work [5–8] has investigated fragmented messages transmission over a *single* disrupted link, modeling packet or file transmission over a wireless link as well as single-hop forwarding of DTN messages. Scenarios where message is delivered over multiple links have not received attention so far.

In this paper, we address the case of message fragmentation over a chain of n disrupted links. This case occurs, e.g., in a static multi-hop wireless network, where link disruptions can be due to interference. We want to estimate the transmission time of a single fragmented message over an empty chain of disrupted links. The message may be rather long. The reason we chose to study this scenario is that it seems to capture one essential aspect of what happens in DTN.

We define a basic model for message transmission over n links in Section 2. The disruptions of communication links in the chain are characterized by i.i.d. ON/OFF periods; the chain is homogeneous in space and time and its links work independently from each other. While the homogeneous chain with identical distributions of disruptions is interesting and mathematically tractable model, none of the practical, actual setups follow exactly all our assumptions. But nevertheless, the model is useful in understanding these practical setups.

In Section 3 we first identify the natural lower and upper bounds on the mean transmission time over n links. Then we derive a generic approximation formula for the mean transmission time. Estimates of the queue sizes in intermediate nodes are needed to compute this formula. In Section 4 we show how to compute these estimates in the cases of uniform and exponential distributions for disruptions. Using these results we can estimate the mean transmission times of fragmented messages in those cases. We stress that while we are using these kind of disruptions to test our formulas, our methods are not restricted to disruptions having exponential or uniform distributions: when the distribution of disruptions has finite mean and variance, the mean transmission times of fragmented messages can be estimated using our methods.

We use relatively simple tools (e.g., one-dimensional random walk and one-step analysis), and try to get as simple as possible formulas. It is possible that even better approximations could be achieved with more refined tools from queueing theory where queues are connected in tandem. Note, however, that queuing theory results are typically about the steady state (long-term behavior) of the system, e.g., the usage of Palm calculus is based on this assumption. These kind of results are not applicable in our case because we are investigating transient behavior. In some sense, the only steady state of our system is the trivial case of the initially empty chain.

To confirm our analysis we have computed the relative error between mean transmission times estimated with our formulas, and the actual transmission times in a simulated environment, where messages are transmitted according to our model over ten disrupted links. From those experiments we conclude that our approximation is suitable for large message sizes, that are at least a few times bigger than what can be typically transmitted within a single contact time; the (relative) accuracy of our estimates increases with the message size, and decreases as we move farther from the source node along the chain.

In Section 5 we derive an alternative approximation method for transmission time that works well for small messages containing only few fragments. The (relative) accuracy of that approximation increases as we move farther from the source node along the chain.

Still in the same Section 5 we derive a recursive lower bound formula for transmission time that works best for small messages divided into many tiny fragments. The tightness of the lower bound decreases when we increase the number of links or the fragmentation unit size.

It can be argued that very tiny messages need not be fragmented at all; if the whole message typically fits into a single contact time there is no point to divide it into pieces. We have a simple approximation formula in Section 3 for the transmission time in this case as well.

We discuss our results in Section 6 and conclude in Section 7. Appendix A contains the justification for the inequality (7) of Section 3.

For ease of reference, we summarize methods for estimating the mean transmission time over a chain of disrupted links in Table 3.

2. System model

The model used to obtain the analytical results is as follows. Network node A sends messages over a chain of n communication links to node B. Nodes are numbered $0, 1, \ldots, n$; node 0 is the sender A and node n is the receiver B.[2] The links change their state between ON and OFF independently from each other in a random manner. This arrangement is illustrated in Fig. 1.

The link speed during the ON state is constant (and the same) for all links.[3] We divide all message sizes by the (constant) link speed, measuring message sizes in seconds. The message size is denoted with x. In particular, for example, when a link is continuously in ON state, a message of length x would be transferred over the link in x seconds. The electromagnetic signal's propagation times and the time it takes to acknowledge transmission over one link are neglected (zero) in our model. We also assume that nobody else (except node A) is sending messages over the chain of communication links. Furthermore, we assume that the link state durations have finite mean and variance.

The sending node A can choose to transmit the message in a single unit, thus requiring sufficiently long contact durations for the whole message to fit. Alternatively, A may split the message into blocks of size f, thus allowing transmission of message fragments consisting of one or more such (equal sized) blocks during shorter

[2] In DTNs, paths of successfully delivered messages are often short because the network diameter is naturally constrained (as, e.g., in deep space networks), or messages do not travel very far in terms of distance and hops (as in mobile opportunistic networks). Therefore, we consider small values of n, say, less than 10, to be more interesting than large ones. But we do not exclude larger numbers of links in what follows.

[3] Please note that a sequence of ON/OFF epochs having different average link speeds during ON epochs, can be transformed into a sequence having constant link speed, that still retains same durations of OFF–ON epoch pairs as the original sequence. The details of this transformation are described in [8].