

Analysis of Bundle Throughput Over LTP

Ricardo Lent, *Senior Member, IEEE*

University of Houston

Houston, Texas 77494

Email: rlent@uh.edu

Abstract—Despite the notable progress that has been achieved on long-haul radio and optical communications, which are expected to improve the data transfer rates of future space exploration missions, both scheduled and random link outages remain possible. To improve reliability, NASA’s delay-tolerant networking (DTN) architecture offers the Bundle Protocol (BP) and the Licklider Transmission Protocol (LTP) that implement bundle custodial transfer and automatic repeat query retransmissions. We analyze the performance of these protocols considering the default flow control method of LTP, which limits the number of concurrent sessions. We derive the combined BP and LTP throughput and provide approximations for model parameters that are difficult to find exactly. Event-driven simulations suggest that our analysis is reasonably accurate under different channel conditions and protocol settings. The proposed model provides insight into the role of the channel conditions, the workload, and the choice of BP/LTP protocol parameters on the performance of a reliable space link.

I. INTRODUCTION

An interplanetary scenario involves extreme communication challenges, such as very long propagation delays, regular link disruptions, asymmetrical performance, and high error rates. To address these issues, the Bundle Protocol (BP)—RFC 5050, which is part of the National Aeronautics and Space Administration (NASA)’s Delay Tolerant Networking architecture, forwards application data encapsulated as bundles over a path. A *bundle* is the name of the protocol data unit. BP creates a store-carry-and-forward overlay to transmit bundles over a path. Custody transfer is used for reliable hop transmissions and the sending node retains the bundle until the next hop on the path receives the transmission successfully.

The Licklider Transmission Protocol (LTP)—RFC 5326, normally operates together with BP to transfer bundles over each link. Depending on the protocol configuration, a bundle may be transmitted as one or more data blocks. Each block consists of a certain number of data segments (DS), whose transmission may occur over several rounds to ensure that the block is correctly delivered at the receiving end. After the first round, subsequent rounds (if needed) attempt to deliver rejected segments due to channel errors using an ARQ (automatic repeat request) procedure. A session is the process of transmitting a block and it may fail after a certain number of attempts. The size of each session is determined by the block size and the segment length, which normally produces equal-length segments (with the possible exception of the last segment). Each segment carries payload and protocol overhead and travels on a separate link-layer frame, such

as the CCSDS/AOS (Consultative Committee for Space Data Systems/Advanced Orbit System) data link.

Knowledge of the expected BP/LTP throughput is essential for the correct planning of the communication capacity of future missions. It is interesting to note that the value of these metrics depends not only on protocol specifications, but also on implementation decisions and parameter settings. The reference LTP implementation available in the Interplanetary Overlay Network (ION) distribution allows multiple concurrent sessions up to a configurable value. The bundle size, block size, and segment payload size are examples of parameters that play major roles on the protocols’ throughput. Nevertheless, limited theoretical results exist to date on the interplay of protocol specifications and implementation choices. The contributions of this work are the:

- 1) Analysis of BP/LTP throughput. Prior works have evaluated BP/LTP throughput experimentally with data delivery times as the main focus. We propose a model for BP/LTP throughput that offers insight into the static (i.e., operator defined) flow control of LTP and the role of the channel, workload, and other protocol parameters.
- 2) Calculation of the maximum (average) number of concurrent sessions. No prior attempts have been reported in the literature regarding this protocol parameter.
- 3) Simulation-based validation of the model. An event-driven simulation that closely follows the channel, LTP, and BP operations provide results that reasonably match model approximations.

II. RELATED WORKS

The performance evaluation of network protocols over space links has been a subject of active research [4], [5]. In relation to LTP, prior works have studied the impact of the convergence layer [13], [15], the aggregation of bundles in blocks [2], [16], and block delivery time [14]. A comparative experimental study of LTP and other space protocols determined that LTP may suffer from lower goodput [8]. Recent works have attempted to improve LTP performance [1], [7], for example through the addition of Reed-Solomon codes [10] or by adjusting packet sizes [3], [9]. A related work has analytically determined the effective capacity of wireless communications [6] and another has experimentally estimated the impact of LTP’s flow control [12] and goodput [11]. This paper addresses the gap in the literature of theoretical methods in BP/LTP throughput analysis.

III. BP/LTP THROUGHPUT

ION's LTP implementation allows the concurrent transmission of more than one block for performance reasons. We focus on 100% red-part blocks (i.e., reliable delivery) and analyze the throughput that can be achieved by the approach. Parameter *number of sessions* (denoted c in our analysis) controls the maximum number of concurrent block transmissions that can be in progress at any given time. In the implementation, the value c is selected manually and the general effect is to implement a static flow control, that is, it allows the operator to regulate the throughput at will. Parameter c is clearly at least 1, but it has no *hard* upper bound. A *soft* upper bound occurs at $c = c^+$, where c^+ represents the lowest number of concurrent sessions on average that are needed to fill the link capacity. However, it is still possible to operate the link with $c > c^+$, although without concrete benefits to the throughput and with a delay penalty caused by the extra buffering times. System constraints, such as the amount of memory available, may also create an upper bound to the value of c . In this study, we define the soft constraint $1 < c \leq c^+$.

To determine the upper limit c^+ , we use the one-way propagation delay d and the observation that each transmission round involves an average of $\bar{n}$ segments:

$$\bar{n} = \frac{1}{\kappa} \sum_{i=1}^{\kappa} np^{i-1} = \frac{n(1-p^{\kappa})}{\kappa(1-p)}$$

given that np^{i-1} segments are transmitted on average at the i -th round when we assume an uncorrelated segment loss probability p . Parameter κ represents the average number of rounds per block that can be approximated with:

$$\kappa \approx \max \left\{ 1, \frac{-\log_p(n) + \gamma}{1-p'} \right\} \quad (1)$$

The formulation results from describing the behavior of LTP rounds as states in an absorbing embedded Markov Chain, where each state j represents the number of segments involved in the i -th round. State 0 is the absorbing state. Observing that the distribution of number of visits to transient state j tends to be proportional to $j^{\log(p)}$, allows to calculate the average number of rounds to absorption as $-\log_p(n) + \gamma$, where $\gamma = 0.5772 \dots$ is the Euler-Mascheroni constant. The value of κ is closely related to the average number of rounds to absorption, but it is low bounded by 1. Also, timeout events contribute to increase κ given that no retransmission can occur during those rounds (except for the checkpoint segment-CP). As a result, the value of κ is corrected by the factor $(1-p')^{-1}$, where p' is the probability of a timeout event:

$$p' = 1 - (1-p_c)(1-p_r) \quad (2)$$

Parameters p_c and p_r are the checkpoint segment (CP) and report segment (RS) loss probabilities respectively, both of which can be obtained using the expression $1 - \sum_{i=0}^{\chi} \binom{L}{i} (1-p_e)^{L-i} p_e^i$, where p_e is the channel BER, L is the frame length and χ is the maximum number of bit errors tolerated for successful frame decoding. We note that p_r requires

knowledge of the return channel BER, which may be different from that of the forward channel. Because the CP segment is always carried by the last DS of each round, its average length may be smaller than a regular DS. We approximate the calculation with $p_c = ep$, for a given constant $0 < e \leq 1$ that is arbitrarily selected.

The upper limit c^+ is the ratio of the segment transport capacity of one round-trip time to the average number of segments per round:

$$c^+ = \frac{2dr/L}{\bar{n}} = \frac{2dr\kappa(1-p)}{b(1+h/m)(1-p^{\kappa})}. \quad (3)$$

Each transmission round lasts for $2d$ plus the transmission time of $\bar{n}$ segments, or equivalently, $\bar{n}/(an)$ bundles. A bundle is rejected with probability $(1-\alpha)^A$ and the channel is available a fraction v of the time. If we assume that the round-trip delay is much larger than the time needed to transmit $\bar{n}$ segments (i.e., beyond cislunar), the bundle throughput (θ_B) can be approximated with:

$$\theta_B \approx \frac{v}{2da\kappa} (1-p^{\kappa})(1-\alpha)^A f(c) \quad (4)$$

where $f(\cdot)$ is a function that describes the throughput increase factor by the concurrency level. The ideal case is $f(c) = c$. Assuming that the application demand is such that a bundle is available for transmission with probability ρ , the BP/LTP throughput (θ) is:

$$\begin{aligned} \theta &= \rho(B - h_B)\theta_B \\ &\approx b\rho v \frac{1-h_B/B}{2da\kappa} (1-\alpha)^A (1-p^{\kappa}) f(c) \end{aligned} \quad (5)$$

It is important to observe that the tail of the distribution of the number of rounds per session can cause $f(\cdot)$ to grow slower than $f(c) = c$. Moreover, increasing c in the vicinity of the theoretical capacity can cause $f(\cdot)$ to grow even slower than linear. These effects are caused by large sessions that transmit less number of segments per round than the average. Further analysis to evaluate the tail of the distribution is still needed. However, we can approximate its impact to the throughput by defining $f(\cdot)$. An empirical expression for $f(\cdot)$ that approximates our simulation results, except for very high BER (e.g., with values worse than 10^{-4}) is:

$$f(c) = p_1 \left(1 - \frac{1}{1 + e^{(p_2+p_4)/p_3}} \right) c \quad (6)$$

where $p_4 = -\log_{10} p_e$. Figure 1 depicts the observed throughput increase factor after increasing the session concurrency level over the single-session transmission case for the validation scenario (see section IV). For each of the BER values that are shown in the figure, which are in the range 10^{-8} to 10^{-3} , the throughput increase factor is shown up to the maximum number of sessions. Using 3, we calculated the maximum number of sessions to be 7, 8, 9, 12, 20, and 89 respectively. The chart shows the case for $m = 200$ except for BER 10^{-3} where $m = 50$. Alongside simulation results, we show a linear regression for each case. Note that expression 6 captures the slope of the linear regressions for each BER

value. We obtained empirically $p_1 = 0.9$, $p_2 = -4$, and $p_3 = 0.5$. With these values, the increase factor introduces a small error for low c , but improves considerably the throughput estimation for high concurrency levels. Another observation is that larger BER values yield to larger c^+ , which may appear contradictory at first. This is due to the fact that high BER sessions normally involve a large number of rounds, each transmitting few segments on average, therefore allowing for higher concurrency.




Fig. 1: Throughput increase factor as a function of the number of concurrent sessions. Larger BER values reduce the benefits of higher concurrency.


IV. ANALYSIS AND MODEL VALIDATION

We compare model outputs to simulation results of BP/LTP. The event driven simulation reproduces the BP/LTP protocol behavior and transmission dynamics on a 100s delay link at 64 Kbps with given BER characteristics and no forward error correction code. We assume that the link is always available ($v = 1$). Statistics of the transmission of 1000 bundles of 2×10^4 bytes were collected to characterize BP/LTP throughput. In the simulation, a bundle fits one block ($a = 1$) and has an overhead of 100 bytes. We selected the segment size as one of the experiment factor. Each segment carries a header of 12 bytes and RS/RA (report acknowledgement) are 20 and 7 bytes long respectively.

The first set of results depicts the observed throughput for different values for the number of sessions c and for a broad range of channel BER values (see Figure 2). The figure includes the reliable link capacity as a reference and results for segment payload lengths of 100 (a) and 2000 bytes (b). Both cases, show a similar trend at different scales with simulation results (marks) following closely analytical results (lines). The usable operational range stops at about 10^{-4} with drastic reductions in the throughput regardless of the packet size. Since the experiment consists in sending files that are


already at the sender side, it corresponds to the case $\rho = 1$ in the theoretical model. It is interesting to note that it was not possible to reach the theoretical link capacity even after increasing the value of c beyond c^+ . The figures include the 95% confidence interval of the mean.

Parameters: $h_b=100, h=7, B=20000, \alpha=1e-06, a=1, d=10, m=100$



(a)

Model: $h_b=100, h=7, B=20000, \alpha=1e-06, a=1, d=10, m=2000$




(b)

Fig. 2: BP/LTP throughput as estimated by the theoretical mode and simulation as a function of the channel BER with segment payload size of 100 and 2000 bytes respectively.


Figure 3 illustrates the impact of the segment payload size to BP/LTP throughput for four BER values: 10^{-3} (a), 10^{-4} (b), 10^{-5} (c), and 10^{-6} (d). BER values smaller than 10^{-6} produce no appreciable difference. The chart indicate a reasonable match between simulation and analytical results.

V. CONCLUSION


With the recent standardization of BP and LTP by the Consultative Committee for Space Data Systems (CCSDS), it is expected that these protocols will continue to play a main role in future space missions. In this work, we have analyzed the BP/LTP throughput over a space link. We developed models that bring insight into the interplay of channel conditions, protocol settings, and workload to communications performance. We analyzed the impact of LTP's default flow control that is implemented by limiting the maximum number

Model: $h_B=100, h=7, B=20000, \alpha=1e-06, a=1, d=10, BER=0.0001$ 

(a)

Model: $h_B=100, h=7, B=20000, \alpha=1e-06, a=1, d=10, BER=1e-05$ 

(b)

Model: $h_B=100, h=7, B=20000, \alpha=1e-06, a=1, d=10, BER=1e-06$ 

(c)

Fig. 3: BP/LTP throughput as estimated by the theoretical mode and simulation as a function of the segment payload size with channel BER of 10^{-4} , 10^{-5} , and 10^{-6} .

of concurrent sessions. A key result of the model is that it allows to analytically determine the maximum number of concurrent sessions for given channel features and conditions. The result is significant because this limiting value is difficult to obtain by other means and only experimental results have been reported so far. Selecting a maximum number of sessions that is too low yields an inefficient use of the channel, whereas

a value that is too high introduces unnecessary buffering and low performance. The proposed model can help to choose the right limit for the number of sessions.

ACKNOWLEDGMENT

The author would like to thank Gilbert Clark for his comments in the preparation of this contribution. This work was supported by an Early Career Faculty grant from NASA's Space Technology Research Grants Program.

REFERENCES

- [1] N. Alessi, S. C. Burleigh, C. Caini, and T. de Cola, "LTP robustness enhancements to cope with high losses on space channels," in *8th Advanced Satellite Multimedia Systems Conference and the 14th Signal Processing for Space Communications Workshop, ASMS/SPSC 2016, Palma de Mallorca, Spain, September 5-7, 2016*, 2016, pp. 1–6.
- [2] N. Bezirgiannidis, S. Burleigh, and V. Tsaoussidis, "Delivery time estimation for space bundles," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 49, no. 3, pp. 1897–1910, July 2013.
- [3] N. Bezirgiannidis and V. Tsaoussidis, "Packet size and DTN transport service: Evaluation on a dtn testbed," in *International Congress on Ultra Modern Telecommunications and Control Systems*, Oct 2010, pp. 1198–1205.
- [4] I. Bisio and M. Marchese, "Analytical expression and performance evaluation of TCP packet loss probability over geostationary satellite," *IEEE Communications Letters*, vol. 8, no. 4, pp. 232–234, April 2004.
- [5] C. Caini, P. Cornice, R. Firrincieli, M. Livini, and D. Lacamera, "Analysis of TCP and DTN retransmission algorithms in presence of channel disruptions," in *2009 First International Conference on Advances in Satellite and Space Communications*, July 2009, pp. 174–179.
- [6] D. A. Carek, "Packet-based protocol efficiency for wireless communications," *JACIC*, vol. 2, pp. 238–251, 2005.
- [7] S. Farrell and V. Cahill, "Evaluating ltp-t: A DTN-friendly transport protocol," in *2007 International Workshop on Satellite and Space Communications*, Sept 2007, pp. 178–181.
- [8] D. Iannicca, A. Hylton, and J. Ishac, *A performance evaluation of NACK-oriented protocols as the foundation of reliable Delay-Tolerant Networking convergence layers*, ser. NASA technical memorandum, 2012. [Online]. Available: <https://books.google.com/books?id=HmOEQAACAAJ>
- [9] H. Lu, F. Jiang, J. Wu, and C. W. Chen, "Performance improvement in DTNs by packet size optimization," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 51, no. 4, pp. 2987–3000, Oct 2015.
- [10] L. Shi, J. Jiao, A. Sabbagh, R. Wang, Q. Yu, J. Hu, H. Wang, S. C. Burleigh, K. Zhao, L. Shi, J. Jiao, A. Sabbagh, R. Wang, Q. Yu, J. Hu, H. Wang, S. C. Burleigh, and K. Zhao, "Integration of reed-solomon codes to licklider transmission protocol (ltp) for space DTN," *IEEE Aerospace and Electronic Systems Magazine*, vol. 32, no. 4, pp. 48–55, April 2017.
- [11] D. Wang, "Performance of licklider transmission protocol (ltp) in leo-satellite communications with link disruptions," in *2016 IEEE 15th International Conference on Cognitive Informatics Cognitive Computing (ICCI*CC)*, Aug 2016, pp. 154–159.
- [12] R. Wang, A. Reshamwala, Q. Zhang, Z. Zhang, Q. Guo, and M. Yang, "The effect of "window size" on throughput performance of DTN in lossy cislunar communications," in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 68–72.
- [13] R. Wang, X. Wu, T. Wang, X. Liu, and L. Zhou, "TCP convergence layer-based operation of DTN for long-delay cislunar communications," *IEEE Systems Journal*, vol. 4, no. 3, pp. 385–395, Sept 2010.
- [14] R. Wang, S. C. Burleigh, P. Parikh, C.-J. Lin, and B. Sun, "Licklider transmission protocol (ltp)-based DTN for cislunar communications," *IEEE/ACM Trans. Netw.*, vol. 19, no. 2, pp. 359–368, Apr. 2011. [Online]. Available: <http://dx.doi.org/10.1109/TNET.2010.2060733>
- [15] R. Wang, B. Modi, Q. Zhang, J. Hou, Q. Guo, and M. Yang, "Use of a hybrid of DTN convergence layer adapters (clas) in interplanetary internet," in *ICC*. IEEE, 2012, pp. 3296–3300.
- [16] Z. Wei, R. Wang, Q. Zhang, and J. Hou, "Aggregation of DTN bundles for channel asymmetric space communications," in *2012 IEEE International Conference on Communications (ICC)*, June 2012, pp. 5205–5209.