Licklider Transmission Protocol (LTP)-based DTN for Long-delay Cislunar Communications

Ruhai Wang[1,2], Paavan Parikh[2], Ramakrishna Bhavanthula[2], Jia Hou[1], Liulei Zhou[1] and Tarik Taleb[3]
[1]School of Electronics and Information Engineering, Soochow University, Suzhou, Jiangsu, P. R. China
[2]Phillip M. Drayer Department of Electrical Engineering, Lamar University, Beaumont, TX, USA
[3]NEC Europe Ltd., Network Research Division, Kurfürsten-Anlage, Heidelberg, Germany
{(rwang, pparikh, rbhavanthu)@lamar.edu, (houjia, zhoull)@suda.edu.cn, tarik.taleb@nw.neclab.eu}

Abstract—Delay/disruption tolerant networking (DTN) technology offers a new solution to highly stressed communications in space environments, especially those with long link delay and frequent link disruptions in deep space missions. To date, little work has been done in evaluating the effectiveness of the available DTN protocols when they are applied to an interplanetary Internet. In this paper, we present an experimental evaluation of the Bundle Protocol (BP) running over various "convergence layer" protocols in a simulated cislunar communications environment characterized by varying degrees of signal propagation delay and data loss. We focus on the Licklider Transmission Protocol (LTP) convergence layer adapter running on top of UDP/IP (i.e., BP/LTPCL/UDP/IP). The performance of BP/LTPCL/UDP/IP in realistic file transfers over a PC-based testbed is compared with that of two other DTN protocol stacks, BP/TCPCL/TCP/IP and BP/UDPCL/UDP/IP. The experiment results show that LTPCL has a significant performance advantage over TCPCL for link delays longer than 4 sec when bit error rate (BER) is 10^{-6} or lower. For a lossy channel with a BER of around 10^{-5}, LTPCL has a significant goodput advantage over TCPCL at all the link delay levels studied, with an advantage of around 3000 bytes/s for delays longer than 1.5 sec.

Index Terms—Satellite communications, interplanetary Internet, DTN, bundle protocol (BP), LTPCL, TCPCL

I. INTRODUCTION

The design of the Internet's transmission control protocol (TCP) [1] is based on some critical assumptions that do not hold in some emerging communication environments like space, resulting in severe performance degradation [2-4]. A new communication architecture called delay tolerant networking (DTN) [5, 6], has been developed to combat the long link delay and frequent link disruptions that generally characterize space communications. To this end, Bundle Protocol (BP) [7], designed to operate at the application layer of the Internet architecture, forms a store-and-forward overlay network to provide custody-based, message-oriented transmission in DTN. The major capabilities of BP include the ability to cope with connectivity interruption and the ability to take advantage of scheduled, predicted, and opportunistic connectivity. BP effects message transmission and reception by invoking the services of an underlying convergence layer protocol (CLP) [7, 8] stack; the nature of this stack may differ in different segments of any end-to-end DTN path, just as the nature of the subnetwork protocol underlying IP may differ in different segments of any end-to-end Internet connection. Currently, the TCP-based CLP (i.e., TCPCL) [8], user datagram protocol (UDP) [9]-based CLP, Saratoga CLP [10], and Licklider Transmission Protocol (LTP) [11−13] CLP are all supported under BP.

TCP and UDP are mainly designed for application in terrestrial networks. In comparison, the newly-developed LTP is designed to operate over point-to-point, long-haul, deep-space radio frequency links or similar links characterized by an extremely long transmission delay and/or frequent interruptions in connectivity. LTP is designed to tolerate these characteristics with no reliance on stability of communication Round Trip Time (RTT). The design of LTP also includes the use of selective Negative Acknowledgments (NAKs) to minimize overhead over asymmetric links, optional accelerated retransmission based on multiple checkpoints per transmission block, and deferred retransmission. However, few studies have evaluated the effectiveness and performance of the DTN protocols, especially LTP, when they are applied to an interplanetary Internet.

In this paper, we present an experimental evaluation of the DTN architecture over a simulated cislunar communication channel in the presence of varying levels of link delay and data loss. The performance of the BP/LTPCL/UDP/IP stack is compared with that of other two DTN protocols, namely BP/TCPCL/TCP/IP and BP/UDPCL/UDP/IP. The experiment was conducted by performing realistic file transfers over a PC-based network test-bed. The intent of the work was, in particular, to investigate the effectiveness of the LTPCL-based DTN protocol in the presence of long link delays and high transmission error rates in a simulated cislunar environment.

In Section II, we describe the experimental configurations and analysis methodologies. We then present the experimental results in Section III. We draw conclusions in Section IV. Finally, some future research work is indicated in Section V.

II. EXPERIMENTAL AND ANALYSIS METHODOLOGIES

A. Experimental Configuration

Table I lists the major protocol configuration and operation parameters of our experiment. Files are transferred via BP/LTPCL over UDP/IP, BP/TCPCL over TCP/IP, and BP/UDPCL over UDP/IP, all over the simulated cislunar communication links. The BP and LTP protocol implementations used are from JPL's Interplanetary Overlay Network (ION) distribution v1.0_r203 [14]. The data rates for both the return and forward channels are 115,200 bit/s. The operating system is Fedora Linux 8 (kernel 2.6.23.1-42).

TABLE I.
EXPERIMENTAL FACTORS AND CONFIGURATION.

Experimental Factors	Settings/Values
Protocol implementations	Interplanetary Overlay Network (ION) v1.0_r203 from NASA JPL, CA
DTN protocol layering	BP/TCPCL/TCP/IP/PPP BP/LTPCL/UDP/IP/PPP BP/UDPCL/UDP/IP/PPP
Operating system	Fedora Linux 8 (kernel 2.6.23.1-42)
Channel rate	115,200 bit/s
BER	0, 10^{-6}, and 10^{-5}
One-way link delay	1280 ms, 1500 ms, 2000 ms, 2500 ms, 3000 ms, 3500 ms, 4000 ms, 4500 ms, and 5000 ms
Experimental file size	1,000,000 bytes

We chose to study three BERs, namely 0, 10^{-6}, and 10^{-5}, to simulate error-free transmission, a moderate error rate, and the maximum acceptable error rate on a channel. The experiment performed with BER=0 is expected to establish the baseline performance of DTN in this environment. The minimal one-way link delay between the Earth and the Moon is around 1.25 s, while the maximal link delay can be as high as 5 s [15]. Within this range, we choose nine different link delays to study (i.e., 1280 ms, 1500 ms, 2000 ms, 2500 ms, 3000 ms, 3500 ms, 4000 ms, 4500 ms and 5000 ms) to evaluate the comprehensive performance of DTN in cislunar communication. The experiment was conducted by transferring a text file of 1Mbyte from the simulated source node on the Moon, TX, through the orbiter, MX, to the destination node on the Earth, RX, of the SCNT test-bed [16, 17]. Each file transfer was performed ten (10) times in each configuration. For a full description of SCNT test-bed development and experimental operation, the interested reader can refer to [16, 17].

B. *Analysis Methodology*

In this work, the experimental results (i.e., most importantly goodput) of each protocol for file transfer are observed from approximately normal populations. We used the *t*-statistic [18] to perform pair-wise goodput performance comparisons among the DTN stacks under study. The *P*-values for three protocol pairs are then compared for each of the three BERs with a widely used reference *P*-value of 0.05, also provided for a comparison reference. Because the alternative hypotheses are created based on a two-tail *t* test, the absolute differences in goodput for each comparison pair are provided in order to signify which protocol has a performance advantage over another and how large the advantage is when they indeed have performance differences. The *P*-value attempts to provide a measure of the strength of the results of a test. It is a measure of how much evidence you have against the null hypothesis. The smaller the *P*-values, the more evidence you have [19]:

- If $P < 0.01$, there is very strong evidence to reject null hypothesis H_0, and accordingly, there is significant difference in goodput between two protocols.
- If $0.01 \leq P < 0.05$, there is moderate evidence to reject null hypothesis H_0, and accordingly, there is reasonable difference in goodput between two protocols.
- If $0.05 \leq P < 0.1$, there is suggestive evidence to reject null hypothesis H_0, and accordingly there is implicative difference in goodput between two protocols.
- If $P > 0.1$, there is no real evidence to reject null hypothesis H_0, and consequently, there is no significant performance difference between two protocols.

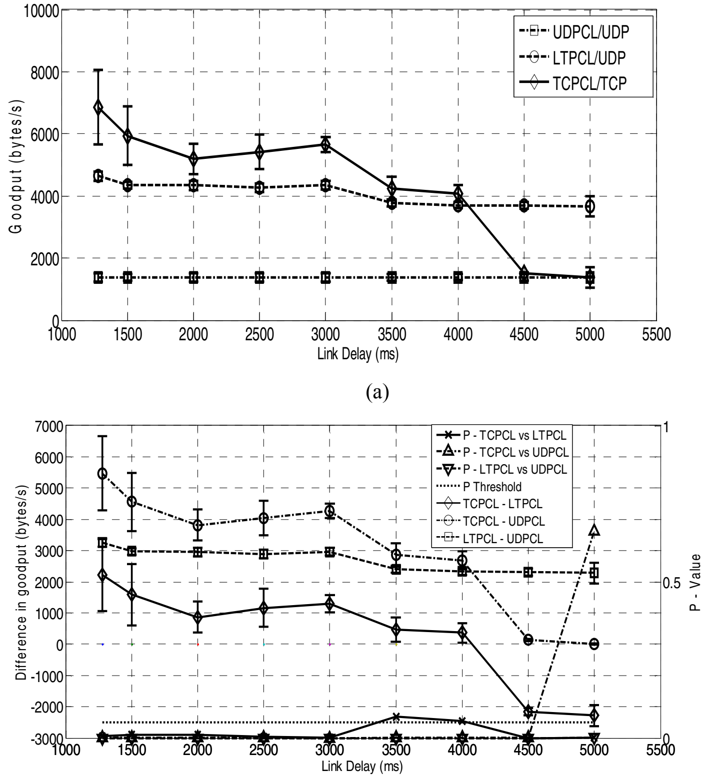

(a)

(b)

Fig. 1. Goodput performance of DTN protocol options in transmission of a 1 Mbyte file over a simulated cislunar channel with constant BER = 0 and varying link delays. (a) Goodput (bytes/s) vs link delay (ms). (b) Difference in goodput (bytes/s) and *P*-values vs link delay.

III. EXPERIMENTAL RESULTS AND DISCUSSIONS

In this section, we discuss the performance evaluation results of the LTP-based DTN protocol stack in comparison with the TCP- and UDP-based DTN protocol stacks. In our discussion, "TCPCL" signifies a transmission with TCPCL as convergence layer running on top of TCP/IP. Similarly, "LTPCL" signifies a transmission with LTPCL as a convergence layer protocol running on top of UDP/IP and "UDPCL" signifies a transmission with UDPCL as a convergence layer protocol running on top of UDP/IP.

A. *Performance over Cislunar Channels with a BER of 0*

Fig. 1 illustrates the goodput performance of three different protocol options when transmitting a 1 Mbyte file over a simulated cislunar channel with constant BER = 0 and varying link delays. Fig. 1 (a) shows the goodput versus delay for the three protocols. Fig. 1 (b) provides the corresponding statistical analysis with difference in goodput and *P*-value versus link delay. The graphic notations in Fig. 1 (b) are explained in Table II. These notations are also applied to the performance analysis of Fig. 2 (b) and 3 (b).

TABLE II.
LIST OF GRAPHICAL NOTATIONS.

Notations	Descriptions
P – TCPCL vs LTPCL	*P*-value for the comparison of goodput performance between TCPCL and LTPCL
P – TCPCL vs UDPCL	*P*-value for the comparison of goodput performance between TCPCL and UDPCL
P – LTPCL vs UDPCL	*P*-value for the comparison of goodput performance between LTPCL and UDPCL
P – Threshold	A widely accepted *P*-value of 0.05 provided for a comparison reference
TCPCL – LTPCL	Averaged goodput differences between TCPCL and LTPCL
TCPCL – UDPCL	Averaged goodput differences between TCPCL and UDPCL
LTPCL – UDPCL	Averaged goodput differences between LTPCL and UDPCL

In Fig. 1 (a), it is obvious that the goodput of the TCPCL protocol is the highest at link delays ranging from 1280 ms to 4000 ms; with further increase in link delay, its goodput drops below that of LTPCL and approximates that of UDPCL. In comparison, LTPCL shows nearly consistent goodput performance, around 4000 bytes/s, as the link delay increases from 1280 ms to 5000 ms; its goodput advantage in comparison with LTPCL and UDPCL is highest when the delay is longer than 4000 ms. All these indicate that link delay has minimum impact on the goodput performance of LTPCL compared to the TCPCL.

We observe that UDPCL shows lower goodput than the two other protocols around 1380 byes/s regardless of link delay. This is not unexpected, since UDP does not include retransmission procedures that would have been affected by variation in round-trip transmission time.

The statistical analysis in Fig. 1 strongly supports the above observations. For the comparison between TCPCL and LTPCL in Fig. 1 (b), the *P*-values, *P*-TCPCL vs LTPCL remain around 0.01 for the link delays from 1280 ms to 3000 ms. This indicates that there is strong evidence to reject H_0 and consequently, there is significant difference in goodput observed between the two protocols for the delays from 1280 ms to 3000 ms. As their differences in goodput (i.e., TCPCL-LTPCL) vary from around 2000 bytes/s to 1000 bytes/s, we conclude that TCPCL has a goodput advantage of around 1000 bytes/s to 2000 bytes/s compared to LTPCL for this range of delays. As the delay increases to 3500 ms and 4000 ms, the *P*-value falls into a new range, 0.05~0.1. This indicates that there is suggestive evidence against H_0 and accordingly, there is minor difference in goodput between two protocols from 3500 ms to 4000 ms. For the delays longer than 4000 ms, the *P*-values are consistently below 0.01. This shows there is strong evidence to reject H_0, which describes the significant goodput difference observed between TCPCL and LTPCL. The goodput differences are around -2200 bytes/s, implying that LTPCL has a goodput advantage over TCP of around 2200 bytes/s for link delays of 4500 ms or longer.

For the comparison between TCPCL and UDPCL, the *P*-value is consistently below 0.01 at all delays except 5000 ms at which it is around 0.67. This indicates there is strong evidence against H_0 for all the delays except 5000 ms at which there is no real evidence to reject H_0. We conclude that TCPCL has a significant goodput advantage, ranging from 100 bytes/s to 5000 bytes/s, over UDPCL for delays shorter than 4500 ms; there is no real performance difference for longer delays.

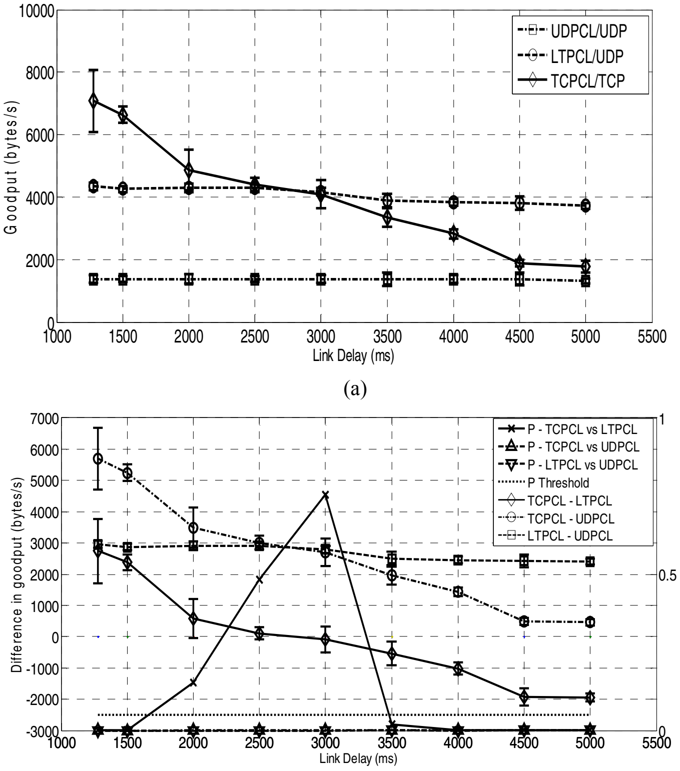

(a)

(b)

Fig. 2. Goodput performance of DTN protocol options in transmission of a 1 Mbyte file over a simulated cislunar channel with constant BER = 10^{-6} and varying link delays. (a) Goodput (bytes/s) vs link delay (ms). (b) Difference in goodput (bytes/s) and *P*-values vs link delay.

For the comparison between LTPCL and UDPCL, *P*-LTPCL vs UDPCL stays consistently in the vicinity of zero at all the delays, indicating that there is very strong evidence to reject H_0 for all delays studied and that the two protocols have significant goodput differences. Their goodput differences, LTPCL-UDPCL, vary from around 3000 bytes/s to 2500 bytes/s. Therefore, we conclude LTPCL has a goodput advantage of around 2000 bytes/s ~ 3000 bytes/s over UDPCL for all the delays ranging from 1280 ms to 5000 ms.

B. Performance over Cislunar Channels with a BER of 10^{-6}

Fig. 2 compares the goodput performance of the three protocol stacks when delivering a 1 Mbyte file over a simulated cislunar channel with constant BER = 10^{-6} and varying link delays. As in the performance comparison at BER = 0, TCPCL has a significant performance advantage over LTPCL for shorter delays, ranging from 1280 ms to 1500 ms. The *P*-TCPCL vs LTPCL is less than 0.01, which is a strong evidence that there is significant goodput difference between the two protocols. With increasing link delay, the advantage of TCPCL disappears and LTPCL starts to show a performance advantage. As observed, the two protocols do not show performance differences for the delays of 2000 ms~3000 ms, as the corresponding *P*-TCPCL vs LTPCL is greater than 0.1. As the delay further increases, LTPCL tends to show a significant performance advantage over TCPCL. This can be seen from the *P*-TCPCL vs LTPCL which is in the vicinity of zero for the delays longer than 3500 ms. The performance advantage of LTPCL is increasing with the increase in link delay, as TCPCL-LTPCL drops from 0 to -2000 byes/s.

Again as in the performance at BER = 0, LTPCL has nearly consistent goodput, around 4000 bytes/s, at all delays, and UDPCL shows consistent goodput, around 1380 bytes/s at all delays, again much lower than the goodput of LTPCL and TCPCL. By the same analytical method as was used at BER = 0, we again conclude that LTPCL has a consistently significant goodput advantage over UDPCL of around 2500 bytes/s to 3000 bytes/s. TCPCL has a significant performance advantage over UDPCL, but the advantage drops significantly from 6000 bytes/s to 500 bytes/s as delay increases from 1280 ms to 5000 ms.

Results at this higher BER differs from those at BER = 0 in that the performance advantage of TCPCL disappears at a much shorter delay level, 2000 ms for BER of 10^{-6} versus 3500 ms for BER of 0. Consequently, LTPCL starts to show a performance advantage much earlier in the experiment, i.e., at much lower levels of link delay.

C. Performance over Cislunar Channels with a BER of 10^{-5}

The performance of these protocols in the transmission of a 1 Mbyte file over a simulated cislunar channel with constant BER = 10^{-5} and varying link delays is presented in Fig. 3. A glance at Fig. 3 (a) shows that LTPCL has a significant performance advantage over both TCPCL and UDPCL at all link delays. This is verified by the comparison of their *P*-values in Fig. 3 (b) where both *P*-TCPCL vs LTPCL and *P*-LTPCL vs UDPCL are below 0.01 implying rejection of the null hypothesis. The advantage of LTPCL over TCPCL is around 1000 bytes/s at delay of 1280 ms and is around 3000 bytes/s for longer delays, and the advantage over UDPCL varies from 3000 bytes/s to 2500 bytes/s. TCPCL has a goodput advantage over UDP of around 2000 byte/s for a short link delay of 1280 ms. As link delay increases above this level, the advantage disappears. When link delay reaches 2500 ms, UDPCL starts to show a performance advantage over TCPCL, and the advantage becomes greater as the delay increases further, approaching 1000 bytes/s at a delay of 5000 ms. The statistical analysis shows that the *P*-TCPCL vs UDPCL is always less than 0.01 except at delay of 1500 ms and 2000 ms where the *P*-values are 0.16 and 0.075, respectively. This means there is no real evidence to reject the null hypothesis and therefore, TCPCL and UDPCL have no significant performance difference with delay around 1500ms. When the delay is increased to 2000 ms, there is suggestive evidence to reject the null hypothesis and accordingly, the two protocols have only implicative goodput difference.

D. Discussions

The operational differences between TCP and LTP explain the discussed performance differences between BP/TCPCL/TCP and BP/LTPCL/UDP over lossy and long-delay cislunar channels.

TCP interprets data loss as an indication of congestion at the receiver (because data loss due to corruption is insignificant in the wired Internet), which causes exponential back-off in the sender's transmission rate. High rates of data loss therefore cause the sender's transmission rate to plummet. LTP, on the other hand, interprets data loss as an indication of corruption and simply retransmits the lost data immediately. This absence of congestion control is a drawback to operation over the Internet but improves performance over a dedicated link such as a cislunar or deep-space communication channel. It should be noted that, in contrast to the immediate retransmission procedures of TCP, LTP may exercise either accelerated retransmission or deferred retransmission as required. These retransmission modes have been shown to outperform immediate retransmission [20].

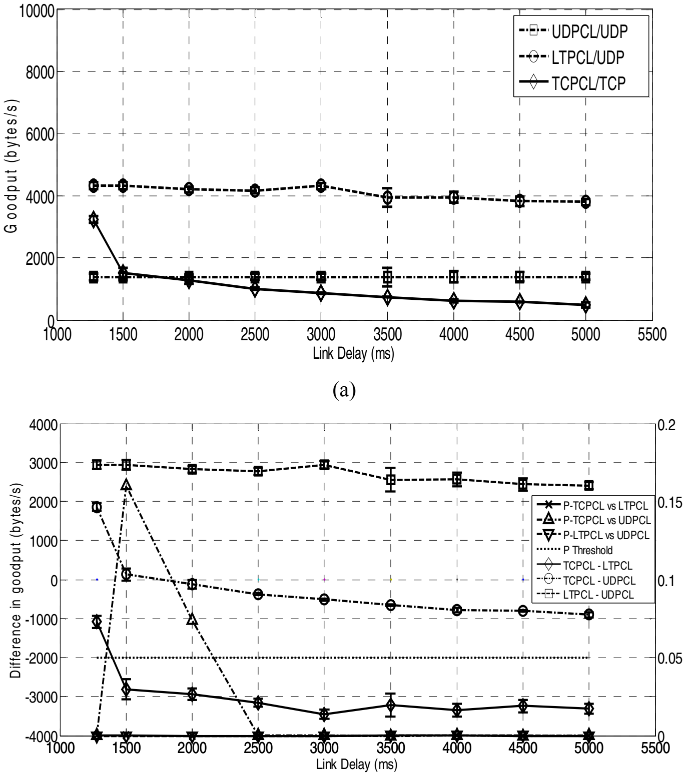

(a)

(b)

Fig. 3. Goodput performance of DTN protocol options in transmission of a 1 Mbyte file over a simulated cislunar channel with constant BER = 10^{-5} and varying link delays. (a) Goodput (bytes/s) vs link delay (ms). (b) Difference in goodput (bytes/s) and *P*-values vs link delay.

Moreover, the TCP sender's reduced transmission rate ramps back up only in response to positive acknowledgments, which are slow to arrive over a long-delay path and it does so in a linear rather than exponential fashion. The result is that restoration of high transmission rate at the sender takes a long time, increasing with the total signal propagation delay over the end-to-end path, and therefore the data rate over the path tends to stay low. LTP suffers no such penalty, because its transmission rate never drops in the first place.

In addition, TCP delivers data only in transmission order. This means that the loss of any single packet will cause a suspension of data delivery for a period of time at least equal to the round-trip time while the request for retransmission is sent and the retransmitted packet is returned. During this time, the rate of data arrival at the receiver will be zero, so again an increase in either bit error rate or end-to-end signal propagation delay will significantly reduce the path data rate.

LTP delivers data in arrival order rather than in transmission order, so the path data rate is unaffected by signal propagation delay or BER.

IV. CONCLUSIONS

In this paper, we have presented an experimental evaluation of a LTP-based DTN over cislunar communication channel characterized by varying levels of link delay and transmission error rate. The major conclusions drawn are as follows:

- One of the strengths of BP is its ability to utilize different convergence-layer protocol stacks that are specifically selected for the different communication environments (e.g., delay, BER, etc.) in which they will be used.
- TCPCL shows significant performance advantage over LTP over a cislunar communication channel with a short delay and a low BER (i.e., 3000 ms or shorter with BER=0 and 1500 ms or shorter with 10^{-6}). As link delay increases, the advantage of TCPCL disappears.
- LTPCL has a significant performance advantage over TCPCL for link delays longer than 4000 ms when bit error rate is 10^{-6} or lower. For a highly lossy communication channel with a BER of around 10^{-5}, LTPCL has a significant goodput advantage over TCPCL at all the link delay levels studied, with an advantage of around 3000 bytes/s for delays longer than 1500 ms.
- LTPCL has a consistently significant goodput advantage over UDPCL, around 2500~3000 bytes/s, at all levels of link delays and BERs.
- Unlike TCPCL, for which the goodput rate is severely affected by an increase in link delay and/or channel noise, link delay and BER have only minor impact on the performance of LTPCL.

To conclude, BP/LTPCL is more suitable than BP/TCPCL and BP/UDPCL for a long-delay and lossy cislunar communication environment.

V. FUTURE WORK

According to the experimental results, TCPCL exhibits performance advantages over less error-prone links while LTPCL has significant advantages over lossy and long-delay space links. We suggest investigating strategies for combining TCP and LTP and evaluating their performance in an integrated space scenario.

According to our research work [21], the traffic shaping mechanism of a rate-based transmission is much more effective than the bursty flow of window-based transmission over space communication channel. There exist protocols such as DS-TP and Saratoga for deep space communications which control data transmissions in a rate-based manner. We will evaluate these new protocols to investigate the performance enhancement of the rate-based mechanism in deep-space communication environment.

ACKNOWLEDGMENT

The author would like to acknowledge Scott C. Burleigh at JPL of California Institute of Technology for his support in configuring ION and conducting the experiment and many significant discussions.

REFERENCES

[1] J. Postel, “Transmission control protocol—Darpa Internet program—Protocol specification,” IETF Request for Comments RFC 793, September 1981.

[2] T. Lakshman and U. Madhow, “The performance of TCP/IP for networks with high bandwidth-delay products and random loss,” IEEE/ACM Transactions on Networking, vol. 5, No. 3, 1997, pp. 336-350.

[3] O. B. Akan, J. Fang, I. F. Akyildiz, "Performance of TCP protocols in deep space communication networks," IEEE Communications Letters, vol. 6, no. 11, November 2002, pp. 478-480.

[4] R. Durst, G. Miller and E. Travis “TCP extensions for space communication,” ACM/Kluwer WINET Journal, vol. 3, no. 5, October 1997, pp. 389-403.

[5] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, R. Durst, K. Scott, and H. Weiss, “Delay-tolerant networking: an approach to interplanetary Internet,” IEEE Communications Magazine, vol. 41, no. 6, June 2003, pp. 128-136.

[6] K. Fall, “A delay-tolerant network architecture for challenged internet,” In Proceedings of SIGCOMM’03, Karlsruhe, Germany.

[7] S. Burleigh and K. Scott, “Bundle protocol specification,” IETF Request for Comments RFC 5050, November 2007, [Online]: http://www.ietf.org/rfc/rfc5050.txt.

[8] M. Demmer and J. Ott, “Delay tolerant networking TCP convergence layer protocol,” IETF DTNRG IRTF Research Group, <draft-irtf-dtnrg-tcp-clayer-02.txt> (Work in Progress), November 2008.

[9] J. Postel, “User datagram protocol,” IETF Request for Comments RFC 768, August 1980.

[10] L. Wood, W. Eddy, W. Ivancic, J. McKim, and C. Jackson, “Saratoga: A delay-tolerant networking convergence layer with efficient link utilization,” In Proceedings of the International Workshop on Space and Satellite Communications (IWSSC), Salzburg, Austria, September 2007.

[11] S. Burleigh, M. Ramadas and S. Farrell, “Licklider Transmission Protocol—Motivation,” IETF Request for Comments RFC 5325, September 2008, [Online]: http://www.ietf.org/rfc/rfc5325.txt?number=5325.

[12] M. Ramadas, S. Burleigh and S. Farrell, “Licklider Transmission Protocol—Specification,” IETF Request for Comments RFC 5326, September 2008, [Online]: http://www.ietf.org/rfc/rfc5326.txt?number=5326.

[13] S. Farrell, M. Ramadas, and S. Burleigh, “Licklider Transmission Protocol—Security Extensions,” IETF Request for Comments RFC 5327, September 2008, [Online]: http://www.ietf.org/rfc/rfc5327.txt?number=5327.

[14] S. Burleigh, “Interplanetary overlay network design and operation V1.8,” JPL D-48259, Jet Propulsion Laboratory, California Institute of Technology, CA, February 2008, [Online]: https://ion.ocp.ohiou.edu/legacy/.

[15] Consultative Committee for Space Data Systems, “Cislunar space internetworking—architecture,” Draft Informational Report, CCSDS 730.1-G-0, Draft Green Book, National Aeronautics and Space Administration, Washington, USA, December 27, 2006.

[16] S. Horan and R. Wang, “Design of a space channel simulator using virtual instrumentation software,” IEEE Transactions on Instrumentation and Measurement, vol. 51, no. 5, October 2002, pp. 912-916.

[17] R. Wang, X. Wu, T. Wang, and T. Taleb, “Experimental evaluation of Delay Tolerant Networking (DTN) protocols for long-delayed cislunar communications,” In Proceedings *of IEEE GLOBECOM 2009*, Nov.-Dec., 2009.

[18] S. Dowdy and S. Wearden, “Statistics for Research,” 2nd Edition, John Wiley & Sons, New York, 1991.

[19] H. Arsham, “Kuiper’s *P*-value as a measuring tool and decision procedure for the goodness-of-fit test,” Journal of Applied Statistics, vol. 15, no. 2, 1988, pp. 131-135.

[20] R. Wang, B. Shrestha, X. Wu, T. Wang, A. Ayyagari, E. Tade, S. Horan, and J. Hou, “Unreliable CCSDS File Delivery Protocol (CFDP) over cislunar communication links,” IEEE Transactions on Aerospace and Electronic Systems, vol. 46, no. 1, January 2010, pp. 147-169.

[21] R. Wang, B. Gutha, and P. K. Rapet, “Window-based and rate-based transmission control mechanisms over space-Internet links,” IEEE Transactions on Aerospace and Electronic Systems, vol. 44, no. 1, January 2008, pp. 157-170.