Memory Dynamics for DTN Protocol in Deep-Space Communications

Jianling Hu, Ruhai Wang*, Xue Sun, Qian Yu
Soochow University and *Lamar University

Zhihua Yang, Qinyu Zhang
Harbin Institute of Technology Shenzhen Graduate School

BACKGROUND

Considering the significant success of the terrestrial Internet, the National Aeronautics and Space Administration (NASA) and other space agencies have been working for a few decades to enable space communications using the Internet-type protocols such as transmission control protocol/Internet protocol (TCP/IP), namely, space internetworking or simply space Internet. It is well known that long propagation delays, intermittent connectivity, heavy channel noise, and asymmetric link rates in space environments all conspire to limit the effectiveness and performance of TCP over space communication channels, especially in deep-space communications [1]. The space Internet in deep-space interplanetary environments is generally named interplanetary Internet (IPN) [2]. Numerous literature surveys [2]–[7] have been done on IPN architectures and protocols.

Many international organizations and researchers have been working on developing new communication protocols and infrastructures for realization of the IPN, and a variety of solutions have been proposed [1], [8]–[13]. Among these developments, delay- and disruption-tolerant networking (DTN) [1] was designed as a new architecture to enable automated network communications despite the long link delay and frequent link disruptions that generally characterize deep-space communications. DTN communications use a bundle protocol (BP) [14] to construct a store-and-forward overlay network that provides custody-based, message-oriented transmission service. To use an underlying convergence-layer protocol stack such as TCP/IP or user datagram protocol/Internet protocol (UDP/IP), BP needs a convergence layer adapter (CLA) [14], [15] deployed between the bundle layer and the transport layer. Currently, the TCP-based CLA (or simply the TCPCL adapter) [15], the UDP-based CLA (or simply the UDPCL adapter) [16], and the Licklider transmission protocol (LTP) [17] CLA (or simply the LTPCL adapter [18]) are the most broadly supported CLAs under BP. Among these CLAs, the newly developed LTPCL adapter [18] is, in particular, designed to operate over point-to-point, long-haul, deep-space radiofrequency links. Some work [19]–[22] has been done in studying the performance of the DTN CLAs in space communication systems. These studies focus on experimental evaluation of file transmission performance and efficiency of the protocols using a test bed.

DTN uses the well-known approach of store and forward with optional custody transfer for which a network node agrees to store a file in memory storage until its successful reception is acknowledged by the next node. While memory consumption for data storage is taken into consideration in protocol design and configuration for the terrestrial Internet, it is a significantly important issue in deep-space communications because of the excessively high cost of space resources.

In a commonly adopted, relay-based IPN communication infrastructure, a relaying spacecraft that takes "custody" of a data unit keeps the data in memory storage until another spacecraft takes over custody or it is received by an Earth ground station. It is important that the memory for data storage of a spacecraft is not occupied by a data unit longer than necessary because of the high cost of memory resources in space. This is critical for file transmission in a deep-space interplanetary infrastructure, which is characterized by frequent and lengthy link disruption and extremely long link delay that likely cause memory storage exhaustion and thus data overflow or losses at a relay spacecraft, resulting in significant performance degradation for the end-to-end data transmission.

While the BP and LTP of DTN are targeted as standard file transfer protocols for deep-space communication networks, it is crucial to characterize its dynamics of memory occupancy and release in file transfer over a range of operational scenarios envisioned for future deep-space missions. However, little work has been done on this subject. In this paper, we present an experimental study of the memory

Authors' current address: J. Hu, R. Wang, X. Sun, Q. Yu, School of Electronics and Information Engineering, Soochow University, People's Republic of China. R. Wang is also with Phillip M. Drayer Department of Electrical Engineering, Lamar University, Beaumont, TX, USA. Z. Yang, Q. Zhang, Department of Electronics and Information Engineering, Harbin Institute of Technology Shenzhen Graduate School, People's Republic of China. Corresponding author is Zhihua Yang, e-mail: (yangzhihua@hitsz.edu.cn). Manuscript SYSAES 2013-0123, DOI. No. 10.1109/MAES.2014.130123, received August 6, 2013, and ready for publication October 21, 2013. Review handled by M. De Sanctis.
0885/8985/14/ $26.00

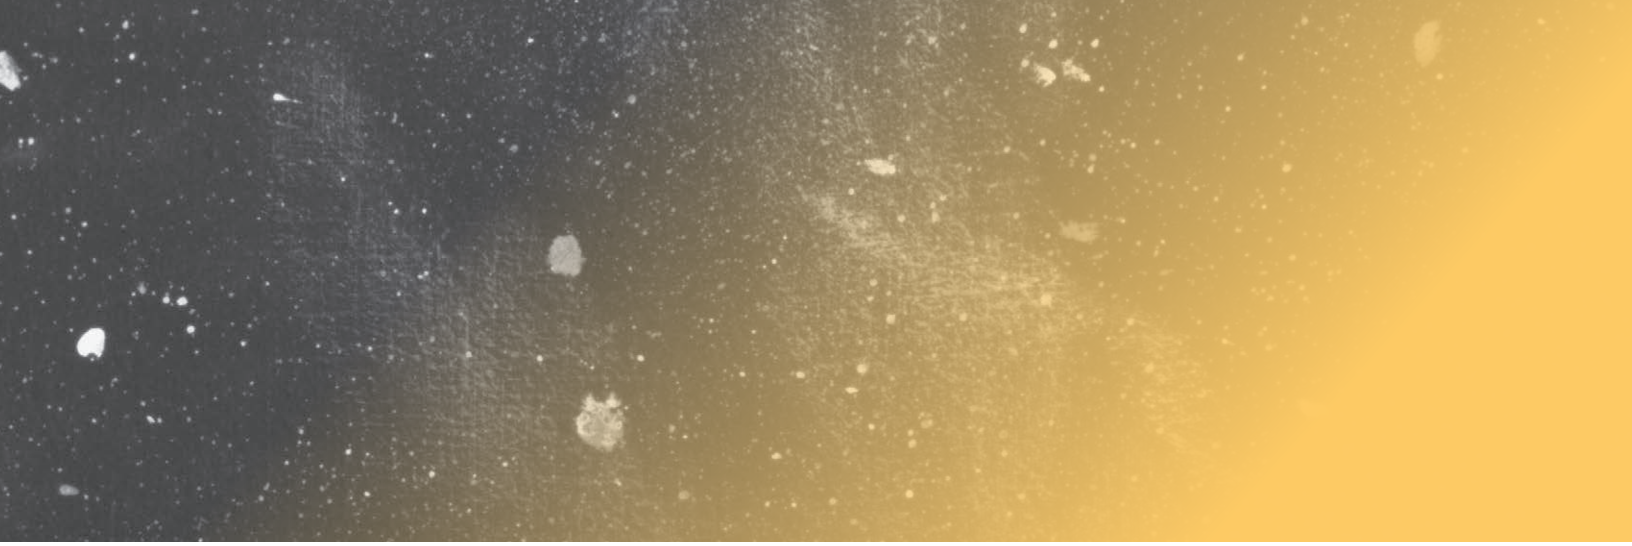

variation dynamics for LTP-based data transmission in a typical relay-based deep-space communication system characterized by an extremely long signal propagation delay, lossy data links, and asymmetric data rates. The goal of this work is to characterize the dynamics of memory occupancy and release imposed by the use of LTP for file transfer in a deep-space mission. Provided that it is hardly possible to study the memory dynamics of LTP for file transfer in a realistic deep-space environment, this work is useful in understanding the use of memory and buffer requirement for LTP in an emulated file transfer scenario that is close to realistic deep space. These are the main contributions of this paper.

OVERVIEW OF LTP-BASED DATA TRANSPORT IN SPACE

As discussed, the LTPCL adapter is the most broadly supported CLA under BP, working with corresponding transport protocols to provide reliable data transport service in the deep-space communications. When LTPCL works with BP, a source file either may be transmitted in a single BP data bundle or may be divided into multiple file data fragments, each of which is transmitted in a separate BP data bundle. For the purposes of the present experiment, each source file was divided into fragments of fixed size for transmission in bundles. BP passes its bundles as "service data units" to LTP for transmission; the service data units are aggregated into LTP client data "blocks" (each block containing one or more complete service data units), and each block is divided into multiple "segments" for transmission [18].

For data transmission using the LTPCL protocol stack, an entire file is transmitted by a sender in a continuous manner in units of data blocks. For the LTP segments of each single block issued by the LTPCL adapter, LTP provides selectable transmission service according to mission requirements and transmission capability, including reliable and unreliable options [18]. Both reliable and unreliable transmission service can be provided for different portions of a single client data block. Figure 1 illustrates an LTP block transmission operation and interactions between sending node and receiving node [23]. As illustrated, a single block consists of a "red" data part for which reliable data transfer is required followed by a "green" data part for which reliable data transfer is unnecessary. The delivery of the "red" part of the block is assured through acknowledgment and retransmission. In comparison, the "green" part is transmitted without any attempt at recovery and completeness if errors occur. The length of either part may be zero; that is, any given block may be designated entirely reliable or entirely unreliable. However, unlike TCP, LTP does not perform any flow or congestion control.

As shown in Figure 1, each data block is fragmented into LTP data segments, according to the underlying link service maximum transmission unit (MTU) size. Some of the segments are flagged as checkpoints (CPs) to check the reception status of the LTP data block. When a CP arrives and all segments of a block are received successfully without error, the segments are reassembled into the original LTP block. The receiver responds to a received CP by returning a report of cumulative reception for the block, termed a report segment (RS), that is, an acknowledgment for the block. In other words, an RS acknowledges a CP, and it serves as either a positive acknowledgment (if all data of this block were successfully received) or a negative acknowledgment (if some data of the block were not successfully received and must be retransmitted). Both RSs and CPs are on timers, and they are retransmitted if not acknowledged. The sender returns a report acknowledgment (RA) to the receiver in response to an RS. The optional transmission of multiple CPs per transmitted block or of interim reports on receptions can provide accelerated retransmission service.

As an illustrative example, in Figure 1, the sender sends the second red segment as a CP to check the reception status of the data block. The CP is acknowledged by the receiver with a RS that the first two red segments were received successfully. Then, an RA is returned to the receiver immediately. Following the transmission of an RA, assume that the last red segment (i.e., the red one marked CP and the end of red part in Figure 1) is sent out but lost; the second green segment is also lost. After an end-of-block message is sent at the end of the file transmission, the lost red segment is retransmitted but the lost green segment is not, because the red data part needs reliable transfer while the green data part does not.

A natural strategy for using LTPCL is for each LTP block to contain exactly one BP bundle (i.e., one bundle per block or, simply, 1 bundle/block). However, as discussed in [21], this strategy results in the transmission of an RS for each bundle, no matter how small the bundle; the volume of acknowledgment traffic can be high, especially if bundles are small. Space communication channels are frequently asymmetric in terms of channel bandwidth: the bandwidth of the uplink (from Earth to the spacecraft) is generally much lower than the bandwidth of the downlink channel (from the spacecraft to Earth). For the highly asymmetric space channels, the limited acknowledgment channel capacity causes delay and possible loss for some RSs and results in goodput performance degradation of data transmission. High channel-rate asymmetry in space constrains aggregating multiple BP bundles into a single block rather than transmitting each bundle in its own block [24]. Because the expanded entire block is still acknowledged by a single acknowledgment (RS) from the receiver, i.e., still following the "one acknowledgment per block" policy, the bundle aggregation may significantly reduce the acknowledgment traffic. As a result, the reduced acknowledgment traffic is likely to be handled effectively by the limited acknowledgment channel rate. However, our prior study [21] shows that bundle aggregation significantly improves LTP performance over highly asymmetric channels, so we exercise that feature in the current work.

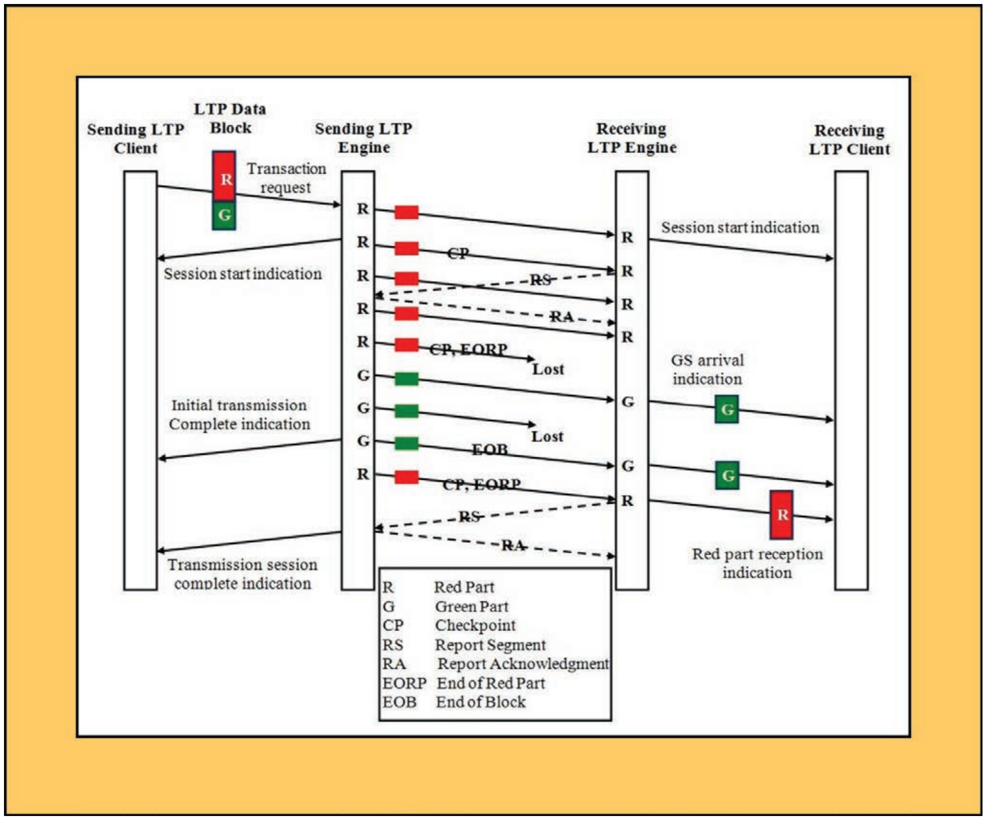

Figure 1.
LTP data block transmission and interactions between sending node and receiving node (courtesy of LTP for CCSDS [23] with changes).

TYPICAL IPN COMMUNICATIONS ARCHITECTURE AND MEMORY DYNAMICS FOR OPERATION OF DTN

Considering the limitations of direct connection time between any two moving planets, a relay strategy has often been taken in space communication to improve the effectiveness of data transmission over a long-haul space link [25]. Recent Mars exploration missions have relied on Mars-orbiting satellites to relay to investigators on Earth the high volume of science data generated by rovers on the Martian surface. These relay operations constitute a functional precursor to the anticipated IPN architecture. In an IPN scenario, data relay could be achieved in two ways: (1) through a planet-orbiting relay spacecraft and (2) through an earth-orbiting spacecraft. In Figure 2, a relay-based IPN (cis-Martian) communication architecture using an orbiting spacecraft to forward data from the Mars to Earth is illustrated [26].

For cis-Martian (i.e., between Mars and Earth) data transmission across the primary data relay path of the IPN architecture shown in Figure 2, the Mars rover, as the source node of data, is the initial custodian for some bundle for which the application has requested custody transfer. It sends the bundle to the Mars relay orbiter as a potential second custodian. After the Mars rover sends the bundle, it also starts a time-to-acknowledgment retransmission timer. If the relay orbiter's BP agent accepts custody of the bundle, it returns a custody acknowledgment to the rover, informing the Mars rover that the data were successfully received and that it accepts the bundle, which is stored for transmission. The rover may then discard the copy of the bundle that it has retained in case retransmission of the bundle is required. If no custody acknowledgment is returned from the relay orbiter before the rover's timer expires, the bundle is retransmitted.

After the bundle is accepted by the Mars relay orbiter, it is immediately forwarded to the third node (which in this case is the destination node on the Earth ground station) if the link between them is available. If no link is available, the bundle is stored (possibly for a long time) in persistent storage at the relay orbiter until the link becomes available. The bundle is eventually delivered to the Earth ground station node, which is acting as the destination node. The bundle is also forwarded to any further custodians in succession in the same way until it is finally delivered to the destination node.

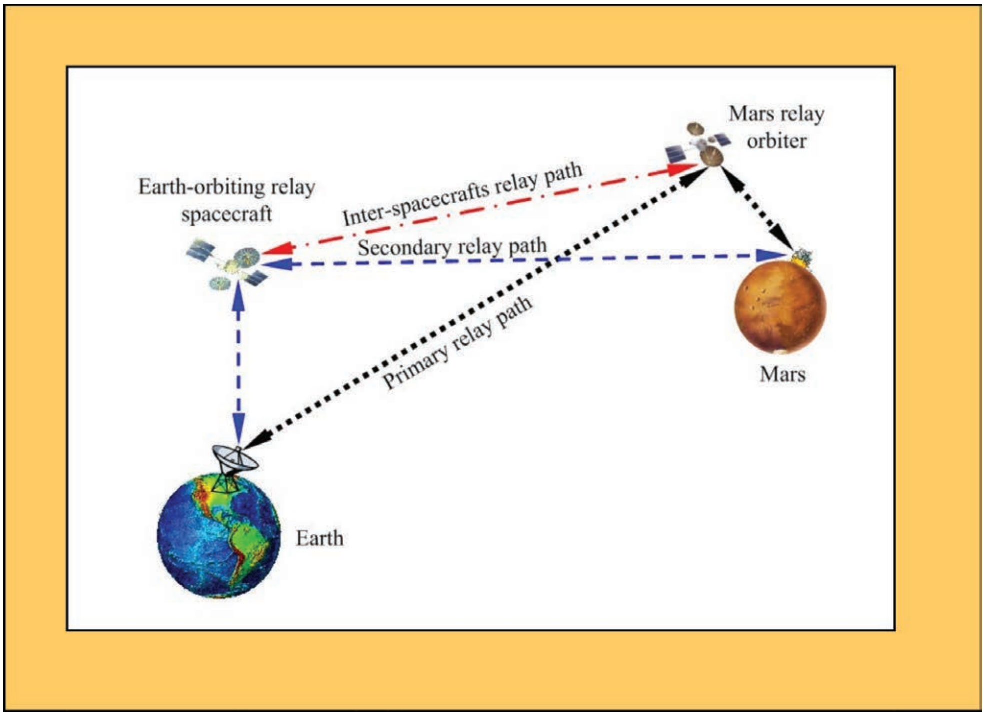

Figure 2.
A typical relay-based deep-space (cis-Martian) IPN communication architecture [26].

The adoption of custody transfer and persistent storage at custodians enhances the likelihood of end-to-end message delivery by advancing responsibility for reliable message delivery toward its final destination. Each custodian (particularly the Mars relay orbiter in this case) keeps a copy of each bundle sent until it receives a custody acknowledgment signal confirming that the bundle has been received successfully at some other node that has taken custody of the bundle. A custody acknowledgment simply signifies that the responsibility for end-to-end reliable delivery of a message has been delegated to another custodian. If the bundles are lost or corrupted when received, they need only be retransmitted from the current custodian, rather than all the way from the source node as in TCP/IP.

Unlike the Internet, in which a router can discard packets when its volatile memory is exhausted, DTN messages are stored in permanent memory storage and cannot be discarded if custody has been accepted. Custodial bundles can only be discarded when either their custody has been accepted by another custodian or their application-specified lifetime has expired. As discussed, it is critical that the memory for data storage is not occupied by a data unit longer than necessary because of the high cost of resources in deep-space communications. Specifically, the mission operation personnel always want the portion of the memory to be released earlier and thus to be made available earlier for other critical space applications. In addition, the frequent and long link disruption and extremely lengthy propagation delay in the interplanetary environment can easily cause the successive hop of a channel (e.g., the link from the Mars relay orbiter to the Earth ground station in Figure 2) to be unavailable for a long time for successive data forwarding. This likely causes memory storage exhaustion and thus data overflow at the Mars relay orbiter in Figure 2. As a result, the lost data have to be retransmitted from the source node to the relay orbiter and possibly have to wait again for the next hop of link to become available. This surely causes transmission performance degradation for the entire data delivery for the mission.

EXPERIMENTAL RESULTS

In this section, we present the evaluation results of the dynamics of memory occupancy and release for LTP data transmission in a deep-space communication system. Due to the similarity of the experimental results, only a few sets of representative results are presented. The experimental setup and configurations are described prior to presentation of the experimental results.

EXPERIMENTAL SETUP AND CONFIGURATION

As introduced, in order to quantify the memory dynamics, realistic file transfer experiments were conducted by running the DTN protocols over a deep-space communication emulation infrastructure. A personal computer (PC)–based space communication and networking test bed (SCNT) was built to implement an emulated deep-space communication infrastructure. Refer to [19, Figure 1] for a block diagram of the SCNT. Previous research [19]–[22] shows that the experimental results obtained from the SCNT are generally considered valid and the test bed can effectively evaluate the performance of a protocol for data-flow file transfer over a space communication channel.

The file transfer experiment with LTPCL was conducted by running BP/LTPCL over UDP/IP on the SCNT, i.e., BP/LTPCL/UDP/IP, which is simply called LTP in this paper. The DTN BP and LTPCL implementations used for our experiments were provided by the Interplanetary Overlay Network (ION) distribution version 3.1.3 [24] developed by the Jet Propulsion Laboratory (JPL), California Institute of Technology (Caltech). For LTPCL, the BP bundle was declared to be 100% red data for reliable transmission of all data bytes. The data of memory occupancy and release

Table 1.

Experimental Factors and Configuration

Experimental Factors	Settings/Values
DTN protocol implementations	ION version 3.1.3 [24] from JPL, Caltech
DTN protocol layering and configurations	BP/LTPCL/UDP/IP with BP's custody transfer option disabled
LTP red/green settings	Bundles are set 100% red data
MTU size	1,500 bytes
LTP segment size	1,400 bytes
Bundle size	1,000 bytes
No. bundles aggregated per block (bundles)	10, 20, 40, 80, 160, and 320
Buffer size	11-Mbyte
Channel ratio (data-to-acknowledgment rate)	1:1 (2:2 Mbit/s)
	100:1 (2 Mbit/s to 20 Kbit/s)
	500:1 (2 Mbit/s to 4 Kbit/s)
BER	0, 10^{-7}, and 10^{-6}
One-way link delay	600 s (10 min)
Experimental file size	10-Mbyte

are collected at the deep-space relay orbiter to measure the memory dynamics imposed by LTP for reliable and complete file delivery in a typical relay-based deep-space communication infrastructure.

Table 1 lists the major configuration parameters of our experiments. The experiments were conducted by transferring a text file of 10-Mbyte using the SCNT, and 16 test runs were performed in each experimental configuration. The source file of 10-Mbyte was divided into fragments of 1,000 bytes that were fixed in size for transmission in bundles. As explained in the overview, BP passes its bundles as service data units to the LTP convergence layer for transmission. The service data units are aggregated into the LTP client data blocks. Six aggregated LTP block sizes are selected for the experiment—10, 20, 40, 80, 160, and 320 bundles/block. These block sizes are selected to study the memory variation dynamics for LTP data transmission with different data transport units. Considering that space communications are generally characterized by asymmetric channel rates, the effect of asymmetric channel rates is also taken into consideration in the file transfer experimental evaluation. Channel ratios of 1:1, 100:1, and 500:1 were selected in the experiment, and their data and acknowledgment channel rates are specified in Table 1. (A channel ratio is defined as the ratio of the data channel rate to the acknowledgment channel rate.)

Three simulated bit error rates (BERs)—0, 10^{-7}, and 10^{-6}—were chosen to simulate three levels of deep-space channel noise. The choice of 0 BER is expected to establish the baseline memory dynamics over a clean or nearly clean channel. The BERs of 10^{-7} and 10^{-6} are adopted to introduce moderate and high data-loss rates over deep-space links, and they are historically typical of the range of noise levels encountered in forward-error–corrected deep-space flight communications. A simulated propagation delay of 10 minutes was selected in our experiment to introduce a deep-space one-way link delay that is common over a cis-Martian channel.

MEMORY DYNAMICS MEASURED FROM EXPERIMENTS

In Figures 3–5, sample experimental results of memory dynamics are illustrated for LTP in transmission of a 10-Mbyte file over an emulated deep-space communication scenario. The evaluation results are only selected for three of the six experimental block sizes—10, 80, and 320 bundles. They are selected to show the memory dynamics at a typical small, midsize, and large block, respectively. The experiments are configured for a scenario with a one-way link delay of 10 minutes, a channel BER of 10^{-6}, and a channel ratio of 500:1 (2 Mbit/s to 4 Kbit/s). The memory dynamics are presented in both the percentile of memory released and the absolute size of memory released versus the file transmission time in seconds. For comparison at all three block sizes shown in Figures 3–5, we observe that the smaller a block size (i.e., with fewer bundles aggregated) is, the earlier the system starts to release the occupied memory. In other words, a larger block (i.e., with more bundles aggregated) generally releases the occupied memory slowly. Numerically, for transmission with a block size of 10 bundles aggregated, shown in Figure 3, around 90% of the occupied memory is released after 20 minutes (i.e., 1,200 seconds), which is the time interval taken by the first transmission effort, given that the configured one-way propagation delay is 10 minutes. For the remaining 10% of the occupied memory, around 8% is released after the second transmission effort, which took another round-trip time of 20 minutes. A few blocks of memory were still occupied after that, and they took another transmission effort to release them, ending up with more than 3,600 seconds required to release all occupied memory. In comparison, for the transmission with a block of 80 bundles, illustrated in Figure 4, around 50% of the occupied memory is released after the first transmission effort, around 40% is released after the second transmission effort, and the rest is released by the third effort.

For the transmission with a large block of 320 bundles aggregated, shown in Figure 5, the occupied memory is released much later. Less than 10% of the memory is released by the first transmission effort, and around 88% of the memory is released after the second effort. In other words, around 90% of the memory is still occupied after the first

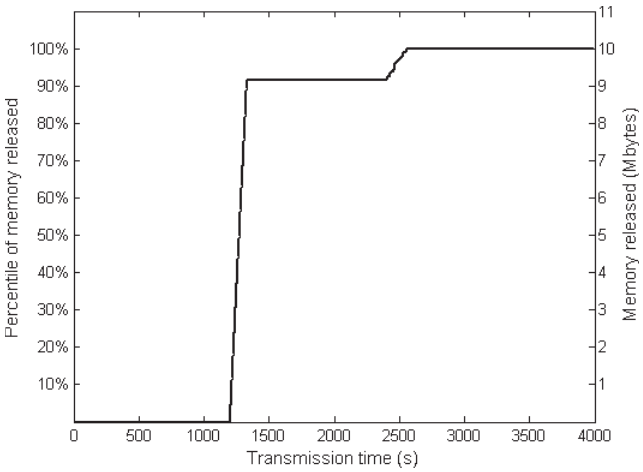

Figure 3.
Memory dynamics of LTP for an LTP block size of 10 bundles (with a bundle size of 1,000 bytes) in transmission of a 10-Mbyte file over a deep-space communication channel with a one-way link delay of 10 minutes, a channel BER of 10^{-6}, and a channel ratio of 500:1.

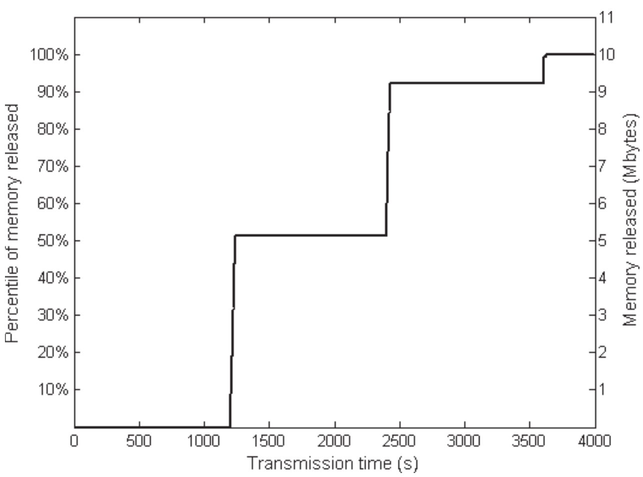

Figure 4.
Memory dynamics of LTP for an LTP block size of 80 bundles in transmission of a 10-Mbyte file over a deep-space communication channel with a one-way link delay of 10 minutes, a channel BER of 10^{-6}, and a channel ratio of 500:1.

transmission effort. From a perspective of memory release performance, this is much slower in comparison to the transmission with a block of 10 bundles, which has only 10% still occupied after the first transmission effort.

The difference in memory release speed among different block sizes can be explained easily. A data sender releases a block of data from the occupied memory only if a positive RS corresponding to this block is successfully received. As described earlier, a data receiver can only issue a positive RS if all segments of a block are completely and successfully received. This means that loss or corruption of a single data segment from a whole block over the data channel causes an entire block of memory to be held by the sender, even when all other segments of the block are successfully received. The sender keeps the memory occupied until the resent (lost or corrupted) segment fills out the gap of the block and an issued positive RS is received.

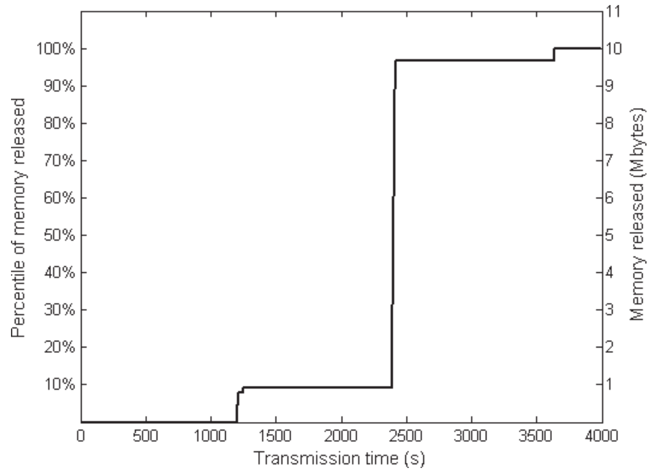

Figure 5.
Memory dynamics of LTP for an LTP block size of 320 bundles in transmission of a 10-Mbyte file over a deep-space communication channel with a one-way link delay of 10 minutes, a channel BER of 10^{-6}, and a channel ratio of 500:1.

Given the same length of a segment (e.g., 1,400 bytes configured in our experiments), a large block size would have a large number of segments for transmission while a small block size has a small number of segments. For example, for the block size of 10 bundles and the block size of 320 bundles experimented with in this work (with the bundle size of 1,000 bytes configured), bundles are transmitted in around 7 and 229 segments, respectively. Statistically, such big difference in the number of segments surely results in a different probability for a successful transmission of a block. This means that on average, a small block requires a much smaller number of transmission efforts (i.e., round-trips) for its successful delivery in comparison to a large block. That implies that the transmission with a small block can receive the corresponding positive RS earlier than the one with a large block; therefore, the sender can release a block of data from the occupied memory faster, because the block is released as soon as the RS is received.

The difference in the memory release speed among different block sizes is more obvious for a comparison among data block sizes. Figure 6 shows a comparison of the memory dynamics among all six experimental data blocks, ranging from 10 to 320 bundles, for the transmission of a 10-Mbyte file at a channel BER of 10^{-6} and a channel ratio of 500:1.

Table 2 provides a comparison of the statistic percentile of memory released during the file transmission after the first round-trip and the second round-trip transmissions for all six experimented block sizes, at all three channel noise levels, and with a channel ratio of 500:1. We see that for all six block sizes, 100% of memory is released at a BER of 0 even after only the first round-trip. But it decreases with an increase in channel BER. For example, at a BER of 10^{-7}, the released percentile is

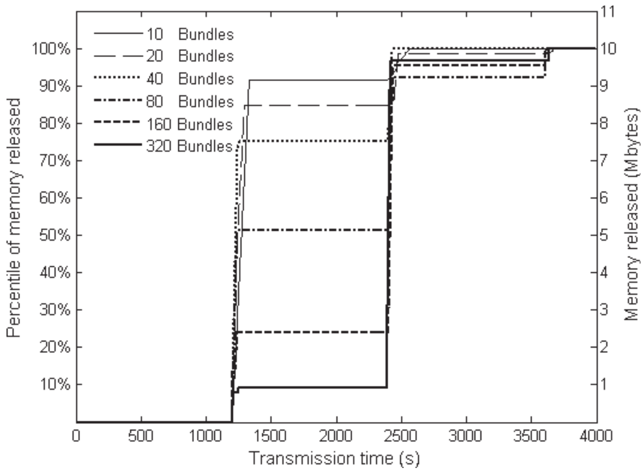

Figure 6.
A comparison of memory dynamics of LTP among all six data blocks ranging from 10 to 320 bundles in transmission of a 10-Mbyte file over a deep-space communication channel with a one-way link delay of 10 minutes, a channel BER of 10^{-6}, and a channel ratio of 500:1.

in a range of 77.83%~99.12%, and no block releases the entire occupied memory after the first round-trip. This is because a higher BER resulted in a larger number of segments corrupted for each block, which generally requires more round-trips to get all bundles retransmitted and successfully delivered to the receiver. The situation is worse for a very large block size at a higher BER; as observed, only 23.89% and 9.21% of the occupied memory are released at a BER of 10^{-6} for blocks of 160 and 320 bundles, respectively. This is because with a high BER, a very large block more likely has more segments that are corrupted and need to be retransmitted, and the occupied memory of a block is released only if all retransmitted segments (and possible re-retransmissions) within the block are successfully received, which results in a positive RS.

The memory release percentile after the third round-trip transmission is not included in Table 2 because all blocks release approximately the entire occupied memory (i.e., 100%) after the third round-trip. This happened because the total number of transmission efforts taken to release the occupied entire memory is three or fewer regardless of the block size, as observed from the illustrated memory dynamics in Figure 6.

The memory release statistic percentile reported in Table 2 is only for a channel ratio of 500:1. However, according to our experimental data, the same variations and trend of the memory dynamics are held for two other channel ratios, 1:1 and 100:1. In other words, given a channel noise level, a particular block size takes the same number of round-trips to release the entire occupied memory, regardless of the channel ratio. In addition, with respect to file transmission performance, our observation shows that for a channel ratio and a given channel noise level, the variations of LTP block sizes have no impact on file delivery goodput.

Table 2.

Comparison of Statistic Percentile of Memory Released During File Transmission for All Experimented Block Sizes at All Three Channel Noise Levels with 500:1 Channel Ratio

	Memory Released During File Transmission		
	BER = 0	BER = 10^{-7}	BER = 10^{-6}
After First Round-Trip			
10 bundles	100	99.12	91.51
20 bundles	100	98.22	84.90
40 bundles	100	96.71	75.05
80 bundles	100	90.91	51.26
160 bundles	100	87.38	23.89
320 bundles	100	77.83	9.21
After Second Round-Trip			
10 bundles	—	100	99.92
20 bundles	—	100	98.71
40 bundles	—	100	99.72
80 bundles	—	100	92.12
160 bundles	—	100	95.34
320 bundles	—	100	96.71

CONCLUSION AND DISCUSSIONS

This paper focuses on characterization of memory consumption and release imposed by the use of LTP for file transmissions over a typical relay-based deep-space communication infrastructure. Based on the file transfer experiment using a PC-based test bed, we conclude that for transmissions over a clean or nearly clean channel (i.e., with a BER around 0), the occupied memory is released quickly regardless of the LTP block sizes—100% released after the first round-trip for all block sizes (having 10 to 320 bundles aggregated, with bundles of 1,000 bytes). This means that the block size has no impact on the memory release speed for transmissions over a clean channel. However, the memory release becomes slower in the presence of channel error, and the speed decreases with an increase in channel BER. In this case, an LTP block size has a substantial impact on memory release: the smaller a block size (i.e., with fewer BP bundles aggregated) is, the earlier and faster a space node releases the occupied memory. This holds true even with variations of channel ratios (1:1~500:1 for experiments in this work) and variations of channel noise level (i.e., BER).

As discussed, the memory components in a space node are generally costly. The drawn conclusions are useful for deep-space mission and protocol designers and mission operation personnel who consider both memory component allocation or usage and transmission performance to be important. For a small block size, the transmission releases the occupied memory much quicker than it does for a large block size; therefore, the portion of the memory released earlier can be made available earlier for other critical space applications, while its transmission performance is as high as that for a large block size.

FUTURE WORK

In this work, we focus on an experimental investigation of memory dynamics for DTN's LTP data transmission in a deep-space communication system. As the planned next-step work, we will perform a similar study in an analytical manner, mainly by building an analytical model to quantify the dynamics of memory occupancy and release imposed by the use of the DTN protocol for reliable and complete file delivery in deep-space missions. By this, we may also compare the experimental results presented in this work with those predicted by the model to verify analytically the conclusions drawn in this work. In addition, an investigation of the memory dynamics of BP with custody transfer would be an interesting and useful companion to the current study.

Some work has been seen in evaluating the transmission performance of DTN protocols in deep-space communications based on simulation or emulation types of file transfer experiments. However, little work has been seen in studying their performance in a theoretical manner. There is an urgent need to build a mathematical model to study the throughput performance of LTP for a better understanding of its operation and transmission efficiency in reliable data delivery in deep space. This should be of significant value to the space DTN research community for adoption of LTP in future deep-space missions and is left as another part of the future work. ◆

ACKNOWLEDGMENTS

The research described in this paper was performed in part at the JPL, Caltech, under a contract with the NASA, and supported in part by the National Natural Science Foundation of China under Grant 61032003. The authors acknowledge Scott C. Burleigh at JPL, Caltech, who has been leading the development of the DTN in space, for the help in implementing the DTN protocols over the test bed. He is particularly appreciated for detailed discussions and careful proofreading of the manuscript that have helped to significantly improve the quality of this paper.

REFERENCES

[1] Burleigh, S., Hooke, A., Torgerson, L., Fall, K., Cerf, V., Durst, R., Scott, K., and Weiss, H. Delay-tolerant networking: An approach to interplanetary Internet. *IEEE Communications Magazine,* Vol. 41, 6 (June 2003), 128–136.

[2] Akyildiz, I. F., Akan, O. B., Chen, C., Fang, J., and Su, W. Interplanetary Internet: State-of-the-art and research challenges. *Computer Networks Journal, Elsevier,* Vol. 43, 2 (Oct. 2003), 75–113.

[3] Wang, R., Taleb, T., Jamalipour, A., and Sun, B. Protocols for reliable data transport in space Internet. *IEEE Communications Surveys and Tutorials,* Vol. 11, 2 (second quarter 2009), 21–32.

[4] Araniti, G., Bisio, I., and De Sanctis, M. Towards the reliable and efficient interplanetary Internet: a survey of possible advanced networking and communications solutions. In *Proceedings of the First International Conference on Advances in Satellite and Space Communications,* Colmar, France, 2009, 30–34.

[5] Bisio, I., Araniti, G., and De Sanctis, M. Interplanetary networks: Architectural analysis, technical challenges and solutions overview. In *Proceedings of the IEEE International Conference on Communications,* Cape Town, South Africa, 2010, 1–5.

[6] Peoples, C., Parr, G., Scotney, B., and Moore, A. A Reconfigurable context-aware protocol stack for interplanetary communication. In *Proceedings of the International Workshop on Satellite and Space Communications,* Salzburg, Austria, Sept. 2007, 163–167.

[7] De Sanctis, M., Rossi, T., Lucente, M., Ruggieri, M., and Mortari, D. Space system architectures for interplanetary Internet. Presented at the IEEE Aerospace Conference, Big Sky, MT, Mar. 2010.

[8] Jiao, J., Guo, Q., and Zhang, Q. Packets interleaving CCSDS file delivery protocol in deep space communication. *IEEE Aerospace and Electronic Systems Magazine,* Vol. 26, 2 (Feb. 2011), 5–11.

[9] de Cola, T., and Marchese, M. Reliable data delivery over deep space networks: Benefits of long erasure codes over ARQ strategies. *IEEE Wireless Communications,* Vol. 17, 2 (Apr. 2010), 57–65.

[10] De Sanctis, M., Rossi, T., Lucente, M., Ruggieri, M., Mortari, D., and Izzo, D. Flower constellation of orbiters for Martian communication. Presented at the IEEE Aerospace Conference, Big Sky, MT, Mar. 2007.

[11] Bisio, I., Araniti, G., and de Sanctis, M. State of the art and innovative communications and networking solutions for a reliable and efficient interplanetary Internet. *International Journal on Advances in Internet Technology,* Vol. 3, 1 (Feb. 2010), 118–127.

[12] Marchese, M., Rossi, M., and Morabito, G. PETRA: Performance enhancing transport architecture for satellite communications. *IEEE Journal on Selected Areas in Communications,* Vol. 22, 2 (Feb. 2004), 320–332.

[13] Caini, C., Cruickshank, H., Farrell, S., and Marchese, M. Delay- and disruption-tolerant networking (DTN): An alternative solution for future satellite networking applications. *Proceedings of the IEEE,* Vol. 99, 11 (July 2011), 1980–1997.

[14] Scott, K., and Burleigh, S. Bundle protocol specification. IETF request for comments 5050, Nov. 2007. [Online] http://www.ietf.org/rfc/rfc5050.txt.

[15] Demmer, M., and Ott, J. Delay tolerant networking TCP convergence layer protocol. IETF DTNRG IRTF Research Group, Nov. 2008. [Online] http://www.ietf.org/internet-drafts/draft-irtf-dtnrg-tcp-clayer-02.txt.

[16] Kruse, K., and Ostermann, S. UDP convergence layers for the DTN bundle and LTP Protocols. IETF DTNRG IRTF Research Group, work in progress, Nov. 2008. [Online] http://tools.ietf.org/html/draft-irtf-dtnrg-udp-clayer-00.

[17] Ramadas, M., Burleigh, S., and Farrell, S. Licklider transmission protocol—Specification. IETF request for comments 5326, Sept. 2008. [Online] http://www.ietf.org/rfc/rfc5326.txt?number=5326.

[18] Burleigh, S. Delay-tolerant networking LTP convergence layer (LTPCL) adapter. IETF Internet-Draft, Apr. 2013. [Online] http://tools.ietf.org/pdf/draft-burleigh-dtnrg-ltpcl-05.pdf.

[19] Wang, R., Burleigh, S., Parik, P., Lin, C.-J., and Sun, B. Licklider transmission protocol (LTP)–based DTN for cislunar communications. *IEEE/ACM Transactions on Networking,* Vol. 19, 2 (Apr. 2011), 359–368.

[20] Wang, R., Wei, Z., Dave, V., Ren, B., Zhang, Q., Hou, J., and Zhou, L. Which DTN CLP is best for long-delay cislunar communications with channel-rate asymmetry? *IEEE Wireless Communications,* Vol. 18, 6 (Dec. 2011), 10–16.

[21] Wang, R., Wei, Z., Zhang, Q., and Hou, J. LTP Aggregation of DTN bundles in space communications. *IEEE Transactions on Aerospace and Electronic Systems,* Vol. 49, 3 (July 2013), 1677–1691.

[22] Yu, Q., Wang, R., Wei, Z., Sun, X., and Hou, J. DTN Licklider transmission protocol (LTP) over asymmetric space channels. *IEEE Aerospace and Electronic Systems Magazine,* Vol. 28, 5 (May 2013), 14–22.

[23] Consultative Committee for Space Data Systems. Licklider transmission protocol (LTP) for CCSDS. Draft recommended standard, CCSDS 734.1-R-2, Draft red book. Washington, DC: National Aeronautics and Space Administration, Feb. 2012.

[24] Burleigh, S. Interplanetary Overlay Network Design and Operation V3.1.3. JPL D-48259, Jet Propulsion Laboratory, California Institute of Technology, CA, October 2013. [Online] http://sourceforge.net/projects/ion-dtn/files/ion-3.1.3.tar.gz/download.

[25] Space Internetworking Strategy Group (SISG). Recommendations on a strategy for space internetworking. Report of the Interagency Operations Advisory Group, July 28, 2008.

[26] Yu, Q., Xue, S., Wang, R., Zhang, Q., Hu, J., and Wei, Z. The effect of DTN custody transfer in deep-space communications. *IEEE Wireless Communications,* Vol. 20, 5 (Oct. 2013), 169–176.