A New Data Priority Paradigm for Deep-Space Communications

Manikantan Ramadas, Shawn Ostermann, and Hans Kruse

Ohio University

`mramadas@irg.cs.ohiou.edu`
`ostermann@eecs.ohiou.edu`
`kruse@ohio.edu`

Abstract. The availability of a powerful data priority mechanism could be extremely useful for communication in the deep-space environment where applications operate under the innate constraints of short communication windows, large values of round-trip times, and high bit error rates. In this paper, we present a new data priority paradigm for applications operating under such environments. This paradigm offers two parameters: Immediacy - to specify the urgency needs of the application data, and Probability of Delivery - to specify the successful-delivery-guarantee requirement sought for the data. The immediacy parameter is valuable in cases of rare and extremely short communication windows to have the most urgent data transmitted and received successfully within the available communication window ahead of other data which, albeit important, can afford to be stored and transmitted later. The probability of delivery parameter serves to distinguish the relative importance/value of the associated data.

1 Introduction

Deep-space presents a challenging environment for data communication. The radio frequency channel predominantly used for communication typically operates under the following constraints:

- *High Signal Propagation Delays* This is almost entirely due to the enormous distances involved between the communicating entities and the relativistic constraint restricting signal transmissions to the speed of light. To cite an example, one-way signal propagation delays for the current Cassini mission to Saturn are in the range of 1 hour and 8 minutes to 1 hour and 24 minutes [9].
- *High Data Corruption Rates* Extremely long distances cause the signals to be received at extremely low strengths at the receiver, and thereby increase the probability of bit-errors in the channel due to random thermal noise errors, burst errors due to solar flares, etc. [10].

- *Disruption Events* Since communicating entities in deep-space tend to be in motion relative to one another, the communication channel between them is prone to disruption. A planetary probe on the surface of Saturn's moon Titan, for example, could experience disruption due to the rotation of Titan on its own axis (when it goes to the night side of Titan), when Titan passes under Saturn's shadow during its revolution around the planet, and when other moons/planets/or the Sun itself block the line of sight to the destination [12].

  Moreover, communicating with an entity in deep-space requires expensive, specialized equipment. The Deep-Space Network system of antennas [14] used for current missions tends to be tightly scheduled for usage by multiple missions and hence the communication window is likely to be restricted from Planet Earth too.
- *Meager, Asymmetric Bandwidth* The bandwidth capacities are asymmetric and fairly limited in the deep-space environment. The uplink channel (Earth to Destination) tends to have much lower bandwidth than the downlink channel (Destination to Earth), typically by one or two orders of magnitude. This is normally by design since the uplink channel is expected to carry mostly light-weight command traffic while all the interesting data collection, analysis, reports etc. are expected in the downlink channel. Therefore the uplink channel tends to be designed as a more reliable, lightweight communication channel. For example, the Cassini spacecraft has an uplink bandwidth of 1 Kbps while the maximum downlink bandwidth is 166 Kbps [9]

2 Motivation

Let us consider the Cassini spacecraft (Figure 1) for example. The Cassini-Huygens mission carries 18 science instruments in total - 12 on the Cassini orbiter and 6 on the Huygens probe to Titan. The science instruments on board Cassini and their estimated data rates [13] are shown in Table 1.

Cassini operates in the X-band (7-12 GHz) radio channel with a downlink data bandwidth (while using its High-Gain Antenna) in the range of 14.2 to 165.9 Kbps, significantly lesser than the total bandwidth required to support all the instruments data simultaneously. It is conceivable that based on the science experimental results, data from the instruments could be of varying degrees of interest and importance during various times in the mission. For example, the Cosmic Dust Analyzer might have an interesting experimental observation that may be of more immediate critical value compared to the last image captured by the wide-angle camera of the Imaging Science Subsystem (which presumably may be stored and transmitted in the next communication window).

Similarly, the Huygens probe carried six science instruments, and its science mission life after being detached from Cassini into the atmosphere of Titan, was only 3 hours (the extreme physical environment within the atmosphere of Titan was not expected to be kind on the communication equipment and batteries). In this short communication window, science data was beamed up to the orbiting

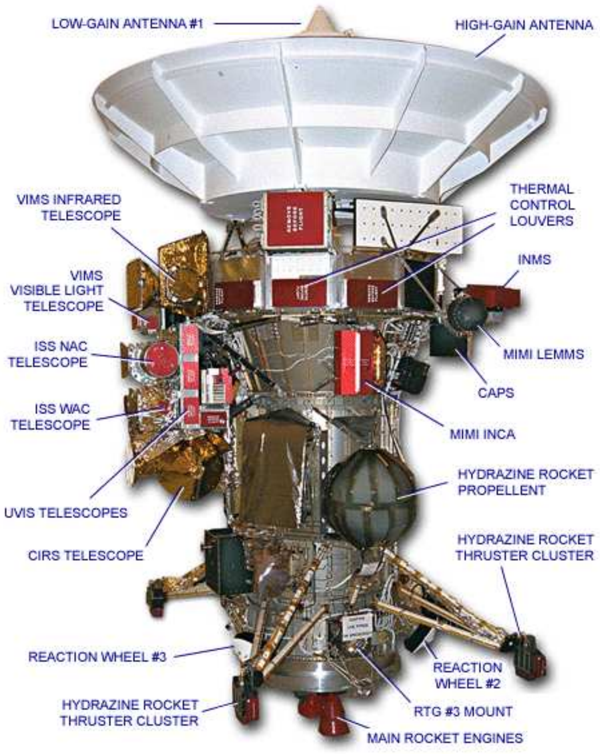

Fig. 1. A View of the Cassini Spacecraft [Image Courtesy : NASA/JPL Caltech]

Table 1. Cassini Science Instrument Data Rates

Science Instrument	Estimated Data Rate Approx. (Kbps)
Cassini Plasma Spectrometer (CAPS)	8.0
Cosmic Dust Analyzer (CDA)	0.5
Composite Infra-Red Spectrometer (CIRS)	6.0
Ion and Neutral Mass Spectrometer (INMS)	1.5
Imaging Science Subsystem (ISS)	365.6
Magnetometer (MAG)	3.6
Magnetospheric Imaging Subsystem (MIMI)	7.0
Cassini RADAR	364.8
Radio and Plasma Wave Science Instrument (RPWS)	0.9
Ultra-Violet Imaging Subsystem (UVIS)	32.1
Visual and Infra-red Mapping Spectrometer (VIMS)	182.8
Total Maximum Downlink Bandwidth	165.9

Cassini spacecraft on S-band radio at a maximum data rate of 16 Kbps. It is obvious that such a short communication window must be optimally used, and preferentially by the applications running on the instruments having the most critical, interesting data.

In addition to the data generated by the science instruments on board a spacecraft, regular housekeeping and telemetry data pertaining to the physical operation and well-being of the spacecraft also needs to be transmitted periodically. It is conceivable that such data is typically not extremely critical, and that healthy operation could be perceived even under the transmission environment where a few telemetry data updates are lost, from the subsequently received updates.

The Licklider Transmission Protocol (LTP) [7] [15] [11] is being designed as a reliable datalink layer protocol optimized to operate over deep-space datalinks. While LTP provides a reliable data-link layer abstraction to the applications, it offers no built-in mechanism to applications for prioritizing their data. On the other hand paper [8], for example, discusses a mechanism for prioritizing images returned by the NASA Mars missions and seems to indicate the need for application level data prioritization requirements. However, there is currently no generic data priority framework for deep-space applications, to the best of our knowledge.

3 Data Priority Paradigm

Our proposed data priority paradigm offers two tunable parameters of control for applications, namely Immediacy and Probability of Delivery.

- *Immediacy* : Immediacy is a notion of how urgently a unit of application data (called a "job" in the subsequent discussion) needs to be received. For example, a message from a science instrument entering a critical state implying status such as "Instrument too hot", "Low battery condition" or commands such as "Stop! Don't go down that cliff" sent from Earth, might need to be reported ahead of all other experimental results / commands. We say that such a job has higher immediacy requirements over the others.
- *Probability of Delivery (PoD)* : In the space environment where there is always a non-trivial bit-error rate on the channel, the communication channel cannot guarantee 100% reliable data delivery. Some application jobs may seek higher PoD guarantees compared to others however. For example, the picture of a microbe found on the surface of Titan may be of much more importance to the mission compared to regular house-keeping telemetry data, and thus may need higher probability of delivery requirements.

A pool of jobs may consist of jobs with varying levels of immediacy and PoD requirements in a continuum. Figure 2 shows four clouds representing just the four corners in the continuum. Our goal is to let the clouds be rather diffuse to let applications specify their priority requirements as desired in the continuum. In the figure, immediacy is specified in units of the maximum number of retransmission cycles within which the corresponding PoD requirement must be guaranteed.

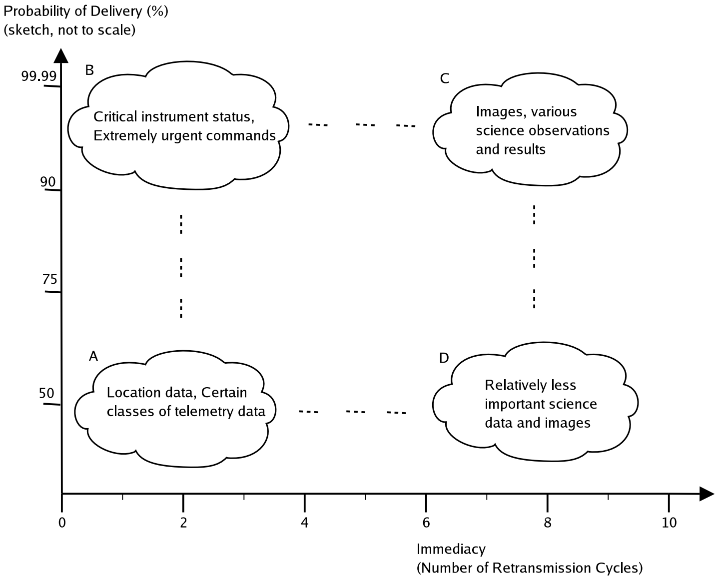

Fig. 2. The Priority Continuum

- Jobs in Cloud A [high immediacy (low actual value), low PoD] could include spacecraft location data and certain classes of telemetry data which must be delivered within a very small timespan or the data would not be valid or valuable any more for the receiver. This includes classes of telemetry data updates that are sent often enough that the loss of a single update is not critical.
- Jobs in Cloud B [high immediacy (low actual value), high PoD] could carry data representing critical instrument status seeking immediate attention ("Instrument too hot" or "Battery draining too fast") or urgent commands sent on the uplink that need to be operated on ahead of all other ongoing jobs ("Stop by that interesting rock for a while"-command being sent to a rover on Mars).
- Jobs in Cloud C [low immediacy (high actual value), high PoD] could include important results of scientific experiments, various interesting images that need to be sent reliably, but not necessarily urgently.
- Jobs in Cloud D [low immediacy (high actual value), low PoD] could include relatively less important scientific data and bulk data that may be sent just on a best-efforts basis. Note that while there typically is a large solid-state

data store available on spacecrafts, its capacity is finite, and an implementation of our priority paradigm would store jobs in Cloud C ahead of jobs in Cloud D, and under severe storage space constraints may even choose not to store Cloud D jobs for transmission.

4 Implementing the Priority Paradigm

Our basic assumption here is that there is always more data to send than the bandwidth capacity of the channel, and that it is important to transmit application jobs as quickly as possible while satisfying their priority needs.

We expect the Bundling protocol [16] to be used by our applications with LTP as the convergence layer. LTP itself may operate over Space Datalink protocols such as the TM [5] [6] / TC [3] [4] protocols or the AOS [2] protocol, standardized by the Consultative Committee for Space Data Systems [1]. We expect the data priority mechanisms outlined here to be implemented in the convergence layer adapter connecting Bundling and LTP.

Here, we consider various mechanisms that could be used to guarantee the priority requirements of application jobs.

- Adapting the error correction mechanisms We might choose to increase the quality of Forward Error Correction (FEC) mechanisms in use for the higher priority jobs and thereby improving the chances of transmitting them successfully, with the additional overhead this would impose.
- Modulating the number of redundant transmissions In a scenario as is typical in current deep-space missions, the FEC mechanisms and frame sizes tend to be fixed for a mission phase. In such a case, the only option we may have to guarantee the job priority requirements might be to transmit the frames redundantly in original transmission. We believe that such a simple mechanism is being used in current missions when the command / data to be sent is extremely critical.
- Adapting the frame size. We might choose to decrease the size of datalink frames for the higher priority application tasks, improving the chances of successful transmission of each frame. This would introduce additional overhead in data transmission as the header data to application data ratio would increase.

Here we discuss a simple scenario where we assume that the probability of bit error on the channel p_b is given to us, and the error distribution is uniform gaussian. We also assume that the FEC mechanisms and frame sizes are fixed and the only parameter we could adapt is the number of redundant transmissions. Let us assume that each frame is of size N bits.

Then we can calculate the probability of frame error p_f to be :

$$p_f = 1 - (1 - p_b)^N$$

For a job J_i with n_i frames, the probability of transmission failure of the job p_{J_i} with single original transmission is :

$$p_{J_i} = 1 - (1 - p_f)^{n_i}$$

We can calculate the probability requirements for the job from r_i to be :

$$R_i = 1 - \frac{r_i}{100}$$

Assuming that the job J_i is a high priority job with an immediacy requirement of 0, we need to guarantee R_i with just the single original transmission. Hence, we then check if the p_{J_i} value we derived was smaller than R_i, in which case we know that the probability requirement is guaranteed. If R_i was smaller than p_{J_i}, we transmit each of the frames γ times redundantly until we find that p_{J_i} meets the R_i requirements.

Note that by transmitting each frame n times, we would arrive at the p_{J_i} value of :

$$p_{J_i} = 1 - (1 - (p_f)^{\gamma})^{n_i}$$

Thus a simple priority scheme could just bump up the γ value for a job until p_{J_i} met the R_i requirement.

The above scenario is fairly simplistic, and might serve to model just the background thermal noise in deep-space datalinks which tends to be fairly uniform gaussian. With a fairly realistic link model characterizing the burst errors on a link, the above calculation would become complex. Such a link model might also require feedback from the receiver on the link performance statistics.

5 Conclusion

We have presented a paradigm for deep-space communications that allows applications to specify their priority requirements in terms of the urgency and importance they associate with their data. Note that this priority paradigm could be valuable in other DTN operating environments too. We also discussed mechanisms to implement the priority paradigm. Our research work is in its early stages at the moment however, and we intend to pursue our research on generating realistic space link models, and scheduling algorithms to transmit a pool of jobs optimally.

6 Acknowledgments

We would like to thank Josh Blanton and other CETS research team members at Ohio University and Scott Burleigh at NASA/JPL for their comments, thoughts, and suggestions related to this work.

References

1. Consultative Committee for Space Data Systems. http://www.ccsds.org.
2. AOS Space Data Link Protocol. Blue Book, Issue 1, September 2003. CCSDS 732.0-B-1.
3. TC Space Data Link Protocol. Blue Book, Issue 1, September 2003. CCSDS 232.0-B-1.
4. TC Synchronization and Channel Coding. Blue Book, Issue 1, September 2003. CCSDS 231.0-B-1.
5. TM Space Data Link Protocol. Blue Book, Issue 1, September 2003. CCSDS 132.0-B-1.
6. TM Synchronization and Channel Coding. Blue Book, Issue 1, September 2003. CCSDS 131.0-B-1.
7. S. Burleigh, M. Ramadas, and S. Farrell. Licklider Transmission Protocol - Motivation, December 2004. Internet Draft (work in progress) draft-irtf-dtnrg-ltp-motivation-00.txt.
8. R. Castano, K. Wagstaff, L. Song, and R. C. Anderson. Validating Rover Image Prioritizations, February 2005.
9. Linda J. Spilker (Editor). Passage to a Ringed World - The Cassini-Huygens Mission to Saturn and Titan, October 1997. NASA SP-533, National Aeronautics and Space Administration, Washington D.C.
10. H. Hemmati et al. Comparative Study of Optical and Radio-Frequency Communication Systems for a Deep-Space Mission, February 1997.
11. S. Farrell, M. Ramadas, and S. Burleigh. Licklider Transmission Protocol - Extensions, December 2004. Internet Draft (work in progress) draft-irtf-dtnrg-ltp-extensions-00.txt.
12. Y. Feria, M. Belongie, T. McPheeters, and H.Tan. Solar Scintillation Effects on Telecommunication Links at Ka-Band and X-Band, May 1997.
13. California Institute of Technology Jet Propulsion Laboratory. Cassini-Huygens Home Page. http://saturn.jpl.nasa.gov/home/index.cfm.
14. California Institute of Technology Jet Propulsion Laboratory. Deep Space Network Home Page. http://deepspace.jpl.nasa.gov/dsn.
15. M. Ramadas, S. Burleigh, and S. Farrell. Licklider Transmission Protocol - Specification, December 2004. Internet Draft (work in progress) draft-irtf-dtnrg-ltp-02.txt.
16. K. Scott and S. Burleigh. Bundle Protocol Specification, September 2004. Internet Draft (work in progress) draft-irtf-dtnrg-bundle-spec-02.txt.