N94-23908

RELIABLE TRANSFER OF DATA FROM GROUND TO SPACE

Fred Brosi

CTA INCORPORATED
Rockville, Maryland, USA

ABSTRACT

This paper describes the problems involved in uplink of data from control centers on the ground to spacecraft, and explores the solutions to those problems, past, present, and future. The evolution of this process, from simple commanding to transfer of large volumes of data and commands is traced. The need for reliable end-to-end protocols for commanding and file transfer is demonstrated, and the shortcomings of both existing telecommand protocols and commercial products to meet this need are discussed. Recent developments in commercial protocols that may be adaptable to the mentioned operations environment are surveyed, and current efforts to develop a suite of protocols for reliable transfer in this environment are presented.

Key Words: Space communications, CCSDS, telecommand, ARQ

1. INTRODUCTION

Transfer of data from control centers on the ground to spacecraft has evolved from sending a few discrete commands, to uploading of command sequences, software modules, and data files. Data rates and volumes have increased by several orders of magnitude over the past 35 years. Total round-trip delay times have increased, due to increased processing and use of LANs both on board and on the ground, and to use of geosynchronous data relay satellites.

In parallel, increasingly sophisticated techniques have been developed to move data from ground to space. Early manual techniques have been replaced by Automatic Retransmission Queueing (ARQ) protocols designed to work in the space link environment. These telecommand protocols meet the basic requirements of complete, in-sequence delivery, but they operate only across the space link, and thus cannot provide reliable transfer end-to-end. They provide only primitive file transfer capabilities, and some require careful monitoring by operators.

Commercial protocols provide many of the end-to-end capabilities needed for space operations, but they were developed for ground-based, multi-user networks, and implementations of these protocols are often optimized for that environment. These protocols lack some of the protocol features needed for efficient use of the scarce resources of the space link (Ref. 1).

The paper concludes by noting recent developments in commercial protocols that may be adaptable to the space operations environment, and the plans of the Consultative Committee for Space Data Systems (CCSDS) to develop Recommendations to meet the space data transport needs of future space missions.

2. EVOLUTION OF COMMANDING

2.1. No Uplink

The first few spacecraft launched in the 1950s had no ability to accept and act on commands from the ground. They simply switched on as they went

into orbit, and transmitted their telemetry signals continuously. Mission objectives were simple, and there was little to control. But before long, mission objectives, and the instruments and spacecraft systems that were sent into space to meet these objectives, became more complicated. We needed to be able to operate the satellite from the ground by sending commands.

2.2. Send & Hope

Early command requirements were modest by today's standards, often limited to turning a few pieces of equipment on or off, or starting or stopping an on-board operation. Sometimes an operating mode was changed. All of these simple operations could be accomplished by use of single, discrete commands. Since each command produced an effect on the spacecraft, the operator could verify that a command had reached the spacecraft and been correctly executed by watching for the desired effect in telemetry on the downlink. The spacecraft had no ability to evaluate the correctness of the sequencing of commands.

Commands sent to a spacecraft do not always arrive intact. Transmission noise, signal fading, or interference could lead to loss of commands. Thus if several commands had to be performed in sequence, the operator would have to verify each command before sending the next one. The communications function of checking for delivery of the commands was performed by the operator and the telemetry system rather than by a communications protocol.

2.3. Send & Count

Later, command procedures were based on a spacecraft command counter, which was incremented each time a new command arrived. The spacecraft also had the ability to reject the remainder of a sequence of commands after detecting an error. The value of the command counter, sent in telemetry, allowed the operator to verify receipt of an entire sequence, or to retransmit any rejected commands. A Barker code at the beginning of each command sequence would unlock the spacecraft, clearing any previous error conditions.

A further refinement of this approach uses sequence numbers in the data units that carry commands. These numbers can be checked on the spacecraft to assure that commands arrive in sequence. This approach is used in the CCSDS COP-1 telecommand Protocol (ref. 2).

2.7. Procedure vs Protocol

As the approaches to commanding described above show, spacecraft telecommand involves a combination of *protocol*—the communications technique used to transfer commands and to control and verify that process, and *procedure*—the actions that the operator must perform to use the protocol and other ground facilities.

Since there are strong drivers to keep space hardware and software simple, commanding protocols have historically been minimal, with the burden falling on operations procedures. But minimal protocols also limit throughput, and as the volume of command and data to be sent from Ground to space increases, more complex protocols are needed to provide this transfer efficiently and reliably.

3. FUNDAMENTAL REQUIREMENTS

The fundamental requirement for a telecommand protocol is complete transfer of a sequence of command data units (CDUs) to the spacecraft in order, and without duplication.

3.1. Completeness

Completeness requires that all CDUs sent from the ground must arrive on board. Either the telecommand protocol or the operations procedure for its use must assure that no CDU goes undelivered. Special features must sometimes be included in the protocol to ensure that the first and last CDU in a sequence is received correctly.

For example, using COP-1, the spacecraft checks sequence numbers to determine if a CDU is missing, but if the last CDU is lost, the spacecraft sees no gap in sequence numbers, and so cannot detect the failure. A timer in the COP-1 sending process is used to initiate retransmission if the last CDU of a sequence goes unacknowledged for too long.

3.2. In Order

In-order delivery requires that if a CDU is lost (e.g., due to an uncorrectable error), any subsequent CDUs in the sequence must not be delivered to their destination on the spacecraft until the missing CDU has been retransmitted and correctly received.

3.3. No Duplication

The requirement for no duplication has a significant impact on the strategy used for retransmitting CDUs when transfer is interrupted. Either the ground must be prevented from sending a duplicate of a command that has already been received by the spacecraft, or the spacecraft must be able to detect and reject duplicates.

4. EVOLVING OPERATIONAL REQUIREMENTS

4.1. Telecommand Context—The CPN

The CCSDS Recommendation for Advanced Orbiting Systems ((ref. 3) defined the concept of the CCSDS Principal Network (CPN), which is the concatenation of on-board, space link, and ground subnetworks. The CPN is the context within which a telecommand protocol operates. As such a protocol is developed, design decisions often depend on the locations of the points within the CPN between which the protocol will operate.

The increasing use of on-board LANs raises the probability of data loss between on-board receipt over the space link and final delivery at the on-board destination. This makes a transport level protocol desirable as a replacement for the link level telecommand protocols now in use.

4.2. High volume Uplink

Larger on-board memories and the need to reduce operations costs through autonomous spacecraft operation is increasing the volume of data transferred from space to ground. This data includes stored command loads, as well as data base updates and software loads. This increase in traffic requires higher rates and higher throughput from telecommand protocols.

4.4. Delay Time

Use of data relay satellites, on-board networks, and more complex ground networks, leads to longer delays. Some multi-spacecraft missions need to pass messages through space-space links to reach end point. These delays, which are orders of magnitude greater than typically seen in ground networks, affect the selection of protocol functions and sizing of protocol parameters. Deep space missions have extremely long delays, and so present a special challenge to designers of space transport protocols.

4.5. Multi-Session Uploads

The need to transfer large data files during short contacts with earth satellites, or over low data rate links to deep space probes leads to a need to spread an upload operation over two or more contacts. Thus a primitive session layer protocol is needed to allow suspension of a transfer operation and restart from a checkpoint.

4.6. Operations Costs

Operating costs are a significant part of the total cost of space missions, and sometimes are the determining factor in limiting the lifetime of a mission. Increasing the level of automation in telecommand protocols simplifies operational procedures, and so decreases the need for manual monitoring and control. As a side benefit, taking the man out of the loop can reduce delay, thus increasing the productivity of a short contact.

5. CURRENT SOLUTIONS FALL SHORT

5.1 COP-1

The current CCSDS Telecommand protocol, COP-1 (Ref. 2), provides all the features needed for reliable transfer at moderate data rates across the Space Link Subnet, but does not provide verification of end to end transfer, and does not support high rate, long delay, or multiple contact transfers.

5.3. SLAP

The Space Link ARQ Procedure (SLAP) defined in (Ref. 3) can transfer data at higher rates and longer delays than COP-1, but it also operates at the link layer, and so is unsuitable for spacecraft with complex on-board networks.

5.2. Send & Dump

The send, dump, compare, and retransmit method used for memory loads in the past has many drawbacks. This method wastes downlink bandwidth and contact time. If an errors are found, either the entire load must be retransmitted, with the same probability of error as in the original transmission, or a new load must be prepared to replace the portions containing the errors. The latter approach is more efficient, but often must be customized to the particular spacecraft, and so is not an approach that promotes standardization.

6. COMMERCIAL PROTOCOLS ARE NOT THE ANSWER

Although the functions needed to support ground to space data transfer are similar to those performed by protocols that operate in ground networks, these commercial, off-the-shelf (COTS) protocols have been designed to work in a different environment from that of spacecraft operations, and efforts to use commercial implementations have, to date, been disappointing. There are two classes of problems: protocol design problems, and protocol implementation problems.

6.1. Protocol Design Problems

6.1.1. Excess Baggage

Most of the protocols that are candidates for space transport protocols have many features and options that are unnecessary or even undesirable. These may include multipoint addressing, extended connection setup procedures, and other features needed in open systems, but which add to the overhead of the protocol data units.

6.1.2. Missing Features

Existing transport protocols, such as TCP (Ref. 4) and TP4 (Ref. 5), lack two important features needed in an efficient spacecraft transport protocol. The first is a means for the receiving

end on the spacecraft to indicate that a retransmission is necessary. Both TCP and TP4 provide for acknowledgement of correctly received data units, but retransmit only when a timer at the sending end expires after a data unit has gone unacknowledged for too long. Given the long, and sometimes variable, delays that are encountered in space operations, this timer must be set to a fairly high value to avoid unnecessary retransmissions. This leads to long periods of unproductive use of the uplink, during which new data units are being sent, only to be rejected by the spacecraft because of an earlier error. This can be avoided by use of an explicit request for retransmission upon detection of an error. Such a feature is provided in the COP-1 and SLAP protocols, and could easily be included in a transport protocol. In (Ref. 6), Zhang concludes that "...we should use external events as a first line of defense against failures, and depend on timers only in cases where external notification has failed."

The second missing feature is an effective means of flow control. Most ARQ protocols use a "credit" or "window" method of flow control, in which the receiving end grants the sending end credit to send a given number of frames before getting an acknowledgement. This works well in the short delay environments in which these protocols were designed to operate. But in the space operations environment, this method can lead to stuttering, in which periods of rapid transmission alternate with idle periods, and under the worst circumstances, result in severely reduced throughput. The causes of this behavior are described in (Ref. 7), along with recommended mechanisms for overcoming it. One mechanism is a rate parameter, which allows the receiving end to regulate the release of CDUs from the sending end to avoid buffer overflow. The SLAP provides this mechanism.

6.2. Implementation Problems

The second class of problems with COTS protocols involves the implementation of the protocols as commercial products. Although these products may provide excellent performance in the ground network environment for which they were designed, experiments show that throughput under nominal space operations conditions can be as low as 5% (Ref. 8).

In (Ref. 8), the factors that contribute to poor performance of a TP4 product under typical space operations conditions are identified. The situation is complex, but is dominated by design decisions that avoid retransmission, under the assumption that data transfer over the connection is to be minimized. When retransmission is needed, the software allows the link to remain idle while waiting to see if a single retransmitted data unit is acknowledged, before trying to send another. In ground networks, where the links are shared by other users, this method controls cost for the user of the protocol, and frees link capacity for other users. But in space operations, the connection cost is dominated by the scarce contact time, and there are often no other users, so this method is counterproductive.

It may be possible to adapt a commercial product to meet space transport protocol requirements, but surgically removing the parts that are not needed may be a more difficult job than building a lean, focused protocol from scratch. Even if the job of extracting the needed subset could be accomplished at a reasonable cost, the interoperability of implementations derived from two different products would be questionable. The cost of space qualifying such software would be at least as great as for an implementation of a customized space transport protocol.

7. FUTURE EFFORTS

Despite the drawbacks to using COTS protocols directly in space operations, these protocols provide a sound starting point for design of leaner, customized space transport protocols. The protocol specifications provide a menu of potential techniques, as well as an established vocabulary for the designers. COTS products can be used for prototyping and testbed evaluation, which helps to identify protocol features that must be modified to achieve needed performance in space operations.

Efforts to develop protocols for very high speed, moderate delay ground networks (Ref. 7) may produce protocols which can be adapted for use in space operations.

The CCSDS has recently begun work to define reliable protocols to meet the evolving needs for end to end transport and file transfer in space operations. This effort will begin with the definition of the functional, performance, and operational requirements for Space Transport and Space File Transfer Protocols.

8. REFERENCES

1. Chitre, Dattakumar M., and Lee, Hsi-Ling 1990. Operation of Higher Layer Data Communication Protocols over Satellite Links. In Proceedings of the IEEE Vol. 78, No. 7, July 1990.

2 CCSDS. 1992. *Space Data System Standards. Telecommand, Part 2 - Data Routing Service: Architectural Specification. CCSDS 202.0-B-2*, CCSDS Secretariat, NASA, Washington, DC.

3 CCSDS. 1989. *Recommendation for Space Data System Standards. Advanced Orbiting Systems, Network and Data Links: Architectural Specification.* CCSDS 701.0-B-1, CCSDS Secretariat, NASA, Washington, DC.

4. Postel, J. 1981. *DoD Standard Transmission Protocol.* ARPA RFC-793. Sept. 1981.

5. ISO/IEC. 1986. *Information Processing Systems - Open Systems Interconnection - Transport Protocol Specification.* International Standard 8073.

6. Zhang, Lixia, 1986. Why TCP Timers Don't Work Well. In *Proc. ACM SIGCOMM Symposium on Communications, Architectures, and Protocols*, Stowe, VT, 1986

7. La Porta, Thomas F., and Schwartz, Mischa 1991. Architectures, Features, and Implementation of High-Speed Transport Protocols. In *IEEE Network Magazine*, May 1991.

8. Durst, Robert C., Evans, Eric L., and Mitchell, Randy C. 1991. The Effect of Long Delay and Transmission Errors on the Performance of TP-4 Implementations. In Proc. IEEE TriComm 91.